



Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles

Report 5 – Phase 2 Implementation of the Concept Demonstrator

*Adam Batten, John Dunlop, Graeme Edwards, Tony Farmer, Bruce Gaffney, Mark Hedley, Nigel Hoschke, Peter Isaacs, Mark Johnson, Chris Lewis, Alex Murdoch, Geoff Poulton, Don Price, Mikhail Prokopenko, Ian Sharp, Andrew Scott, Philip Valencia, Peter Wang, and Denis Whitnall
CSIRO Telecommunications & Industrial Physics, Lindfield, New South Wales, Australia*

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.

- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include creating custom thesauri, building customized databases, and organizing and publishing research results.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at <http://www.sti.nasa.gov>
- E-mail your question via the Internet to help@sti.nasa.gov
- Fax your question to the NASA STI Help Desk at 443-757-5803
- Phone the NASA STI Help Desk at 443-757-5802
- Write to:
NASA STI Help Desk
NASA Center for AeroSpace Information
7115 Standard Drive
Hanover, MD 21076-1320



Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles

Report 5 – Phase 2 Implementation of the Concept Demonstrator

*Adam Batten, John Dunlop, Graeme Edwards, Tony Farmer, Bruce Gaffney, Mark Hedley, Nigel Hoschke, Peter Isaacs, Mark Johnson, Chris Lewis, Alex Murdoch, Geoff Poulton, Don Price, Mikhail Prokopenko, Ian Sharp, Andrew Scott, Philip Valencia, Peter Wang, and Denis Whitnall
CSIRO Telecommunications & Industrial Physics, Lindfield, New South Wales, Australia*

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Purchase Order L-71308D

September 2009

The use of trademarks or names of manufacturers in this report is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products or manufacturers by the National Aeronautics and Space Administration.

Available from:

NASA Center for AeroSpace Information (CASI)
7115 Standard Drive
Hanover, MD 21076-1320
443-757-5802

Table of Contents

| | |
|---|-----------|
| 1. Introduction | 4 |
| 2. Overview of the System Hardware | 7 |
| 2.1 Introduction | 7 |
| 2.2 Physical Structure | 7 |
| 2.3 Modules and Layers | 8 |
| 3. Sensors, Sensing and Mechanical Mounting..... | 10 |
| 3.1 Simulation of fast-particle impacts..... | 10 |
| 3.2 Sensors..... | 10 |
| 3.3 Aluminium panels | 11 |
| 3.4 Mechanical fixtures | 11 |
| 3.5 The Concept Demonstrator | 18 |
| 3.6 Determination of impact location..... | 18 |
| 4. Hardware and Software Implementation of the Physical Cells..... | 22 |
| 4.1 Data Acquisition Sub-Module..... | 22 |
| 4.2 Network Application Sub-Module | 23 |
| 4.3 System Integration..... | 24 |
| 4.4 Software and Configuration Updates | 24 |
| 4.5 Network Communications and Protocols..... | 26 |
| 5. System Operation Workstation and Visualization..... | 32 |
| 5.1 Visualization software..... | 32 |
| 6. Intelligent Distributed Processing..... | 36 |
| 7. The Demonstration System..... | 39 |
| 8. Summary and Conclusions | 42 |
| 9. References | 43 |
| Appendix A: Location of Impacts..... | 44 |
| A.1 Introduction | 44 |
| A.2 Analysis | 45 |
| A.3 Conclusions | 53 |
| Appendix B: Publications, Reports, Conference Papers and Articles Submitted..... | 54 |

1. Introduction

NASA's goal of ageless or robust aerospace vehicles (RAVs) requires the development of vehicles that are capable of structural self-assessment and repair. This project is concerned primarily with the self-assessment, or health monitoring, functions at this stage, but it will progress towards the incorporation of damage prognostics to enable intelligent, forward-looking damage mitigation decisions to be made. The health monitoring functions can be divided between those carried out by distributed sensors and intelligent processing and communication on the skin or within the structure, and those that could be more effectively provided by autonomous robotic NDE agents which could be deployed to monitor damage or integrity of the vehicle structure.

Critical to the success of all structural health monitoring programs is the development of appropriate technologies for non-destructive evaluation of structures, and the development of strategies and technologies for processing NDE data, storage and communication of NDE information, and analysis of NDE data with the capability for intelligent decision-making.

Previous work in this project developed and examined concepts for integrated smart sensing and communication systems that could form the distributed sensing function of a smart vehicle. This work was outlined in four earlier Reports (CTIP, 2001; CTIP, 2002; CTIP, 2003a; CTIP, 2003b). The third of these Reports (CTIP, 2003a) contained a proposal for the design of a Concept Demonstrator and Experimental Test-bed, a combined hardware and software system capable of demonstrating principles of an intelligent vehicle health monitoring system. The fourth Report (CTIP, 2003b) is a brief outline of the first stage of development of this Concept Demonstrator (CD), and some of its features and capabilities.

The purpose of the CD is to demonstrate concepts for an intelligent vehicle health monitoring system within a relatively simple environment. It was decided that, at least in the first instance, the only threats to the structure that would be considered would be from impacts by projectiles that, for a vehicle in space, might be micrometeoroids or space debris. A summary of threats associated with particle impacts, compiled from a number of sources, appears in Report 2 (CTIP (2002), Section 5.3 and Appendix A5.1). *Thus, the aim of the initial design of the CD is only to enable the detection, location and evaluation of the effects of particle impacts.* It is expected that other threats and damage processes will be addressed in later stages of the project.

One of the most important principles of the design, and one that distinguishes a health monitoring system from other intelligent sensing systems (or 'smart spaces'), is the requirement for continued functionality in the presence of damage. A system designed to detect and evaluate damage, and ultimately initiate repair and/or other remedial action, clearly must continue to operate effectively when damage occurs, whether it is damage to the vehicle structure, damage to the sensing system, or damage to both.

The requirements for robustness, reliability and scalability in such a system led to the adoption of a complex multi-agent system approach to the development of intelligent diagnosis, prognosis, decision-making and response. This approach has been discussed in our earlier Reports, and examples given of algorithms and the resulting emergent behaviours that can be produced in such a system. One of the major research challenges in this area is the design of agents and algorithms that produce desired emergent responses. Agent-based hybrid systems

will be developed to provide solutions in the shorter term, but a full complex multi-agent approach is believed to be the ultimate solution, and this is the path that is being pursued in this work. The CD will provide an extremely valuable test-bed for experimentation with and evaluation of multi-agent algorithms.

It should be emphasized that the CD is NOT intended to be a prototype of a practical vehicle health monitoring system. The design of practical systems will always be subject to constraints of, *inter alia*, cost, weight, power consumption, EMI with other vehicle functions, etc., but effective design will require an understanding of the performance trade-offs involved in applying the constraints to varying degrees. The CD design took no account of any of these constraints except cost, which is always limited by available resources. The present work aims to develop the principles of high performance structural health monitoring systems: it is hoped that one of its side benefits will be a greater appreciation of the compromises inherent in applying various constraints to the design of practical systems.

This Report is intended to be an up-date of Report 4 (CTIP, 2003b), and it will outline the developments that have taken place over the past six months. Details of the purpose and design of the CD were presented in Report 3 (CTIP, 2003a), and they will not be repeated here. However, there have been some modifications to the original plans, driven largely by the magnitude of the engineering task involved in the development of the CD hardware and communications software, and by delays in deliveries of electronic components and manufacture of modules.

The Statement of Work for the current contract (NASA Langley PO Number L-71308D, April 28, 2003, Base Option) specifies the following.

Further development of the sensing and intelligent decision-making capabilities of the Concept Demonstrator. In particular:

- Implement up-dated versions of network and acquisition hardware, and complete fully functional Concept Demonstrator.
- Demonstrate system awareness of locations and severities of multiple impacts.
- Develop robust characteristics of the intelligent system, to enable dynamic reconfiguration in the presence of damage to one or more cells.

These requirements have not been fully met, as will be detailed below, but significant progress has been made. A second (and enhanced) version of the Network Application Layer (NAL) hardware has been completed, and a complete network of 192 cells installed on the Concept Demonstrator. A first version of the Data Acquisition Layer (DAL) hardware has been designed, prototyped, tested (and subsequently modified) and a small number fabricated. It was not possible to produce a full complement of these boards because production problems at Texas Instruments delayed delivery of the TMS320F2810 processor chips, and because of an equipment failure at the board manufacturing contractor's plant. These issues will be covered in greater detail later in the Report. As a result of this the Concept Demonstrator is not yet fully functional, as only a small number of its sensing agents have sensing capability. Nevertheless, the system has complete network functionality, and a substantial amount of its functionality can be demonstrated.

The requirements for the next phase, Option 1 of this contract, to be completed by 30 September 2004, are as follows.

- Investigate the use and implementation of dynamically determined hierarchies among the cells, with different modes of communication (e.g. wireless, IR) being used for communication among different levels of the hierarchy.
- Implementation of impact damage evaluation, using additional sensors or using the existing sensors differently to obtain additional information about damage severity.

The first of these points relates to the use of hierarchies of cells in the diagnosis and prognosis of damage. Whether it will be appropriate and/or advantageous to use different communication modalities within different levels of a hierarchy will be the subject of investigation. The second point relates to the development of secondary sensing techniques, probably involving active sensing, for the evaluation and diagnosis of damage that has been detected using the primary (passive) sensors.

The remainder of this Report is set out as follows, and it should be borne in mind that it is intended as an up-date to Report 4 (CTIP, 2003b), which described the initial phase of the CD implementation. The next Section contains a general overview of the CD system hardware in its present form. This is followed by a description, in Section 3, of the sensors, and the mechanical mounting of the sensors and electronic modules to the aluminium ‘skin’ panels. Then follow Sections that outline recent developments to the physical cells’ hardware, software and inter-cell communication protocols (Section 4), and the system control workstation and visualizer (Section 5). Recent developments of intelligent agent algorithms and software are briefly outlined in Section 6, followed by a description of the demonstration system as it was in April 2004 (Section 7), and a brief summary and conclusions. There are also two Appendices: Appendix A provides an analysis of the accuracy of impact location with the sensor arrangement employed in the CD, and Appendix B lists the written research output (published papers, reports, conference papers and submitted articles) of the project group since the work commenced in November 2001.

2. Overview of the System Hardware

2.1 Introduction

The initial goal of the present CD system is the detection and characterization of high-velocity impacts, such as might be caused by micrometeoroids in space. Passive piezo-polymer sensors 2.5 mm in diameter have been bonded to the inside surface of the vehicle ‘skin’, with an average distance between sensors of about 60 mm, to detect the elastic waves generated in the structure by impacts. The sensors, processing and communications hardware are constructed as modules, each containing the sensors on roughly 100 mm × 100 mm areas of the aluminium sheet skin, with the electronics being mounted in layers directly behind the sensors. This unit is termed a ‘cell’. A schematic diagram of the system is given in Figure 1.




Figure 1: Schematic diagram of the test-bed, containing the physical structure and cells, simulated cells, and test-bed control workstation. Each side of the Concept Demonstrator hexagonal structure is divided into 4×2 panels, each of which contains 2×2 cells, which are the basic modular units of the system.

2.2 Physical Structure

The structure of the Concept Demonstrator is a hexagonal prism, with each rectangular face being 440 mm × 800 mm (note that the width of each face of the prism has been increased during the current reporting period). A modular aluminium frame is covered by 220 mm x 200 mm, 1-mm thick aluminium panels that form the outer skin of the structure. Each such panel contains four cells, and each of the six sides of the prism contains eight of these panels, so the entire Concept Demonstrator contains 48 panels and 192 cells (see Figure 1). As each cell contains four sensors there are a total of 768 piezo-polymer sensors in the initial system.

This structure has been placed on a platform that enables it to be rotated, elevated and translated, which simplifies testing, as it is easier to move the structure than the source of impacting particles (to be described later). As the design has been modularized, the same hardware could be readily applied to a different structure if desired.

Although the initial system uses a single type of sensor to measure a single threat, namely high-velocity impacts, the system has been designed so that in the near future the CD may be extended to use multiple sensor types to detect and characterize a range of threats.

2.3 Modules and Layers

A modular approach to the sensing and electronics was adopted to enhance flexibility, re-configurability and ease of manufacture (see Hedley et al., 2003 for more details). Each module, or cell, contains sensing elements, signal processing (analog and digital) and communications, and there is a logical layering of these functions, as follows.

1. Sensors – initially piezoelectric polymer sensors only are attached to the aluminium skin.
2. Analog signal processing – this includes amplification, filtering and other processing of the sensor signals required prior to digitization of the signals.
3. Sampling and data pre-processing – this includes digitization, calibration correction, data reduction and other processing that can be performed using only the local signals.
4. Data analysis – at this level data is processed using information from local sensors and neighboring modules. This is the layer at which the agent-based algorithms will be implemented to provide the network intelligence.
5. Inter-module communication – this layer comprises the software stack and the physical links to provide communication between modules.

These layers are divided into two groups, the Data Acquisition Layer (DAL), which consists of layers 1 to 3, and the Network Application Layer (NAL), which consists of layers 4 and 5. These are implemented as separate physical sub-modules, called the Data Acquisition Sub-module (DAS) and Network Application Sub-module (NAS) respectively. This is illustrated in Figure 2. This separation allows replacement of one type of sensor, and its associated electronics, with another sensor sub-module, without changing the main processing and communications hardware, hence simplifying the use of a range of sensor types in the test-bed. There are separate digital signal processors (DSPs) in each sub-module, and these are called the Data Acquisition Processor (DAP) and Network Application Processor (NAP) respectively in each layer.

Each DAL sub-module communicates with one NAL sub-module, and network communication is provided by each NAL sub-module communicating with its four nearest neighbours. This provides a highly robust mesh network structure that will maintain connectivity even when a significant number of cells or communications links are damaged.




Figure 2: Layer structure and communications of the cells. Each cell contains a Data Acquisition Sub-module (DAL) and a Network Application Sub-module (NAL), with communications links between the Network Application Sub-modules forming a mesh network structure

3. Sensors, Sensing and Mechanical Mounting

3.1 Simulation of fast-particle impacts

Report 4 (CTIP, 2003b) showed that laser pulses from a Q-switched Nd:YAG laser generate elastic waves in aluminium plates that exhibit all the characteristics of waves generated by high-speed particle impacts. In this phase of the project all the testing has been done with laser pulses rather than particle impacts. This provided a much more convenient and faster way of testing various sensor configurations and the signal processing electronics, and ultimately the operation of the whole Concept Demonstrator itself. Pulses at repetition rates of up to 20 Hz are possible, and the energy of the pulses can be changed reliably over a range of nearly three orders of magnitude. This flexibility of the laser source, compared to the few shots per day allowed by a gas gun firing small particles, made the laser the obvious choice for testing during this stage of the project.

3.2 Sensors

Piezoelectric sensors, consisting of a 110 μm -thick film of PVDF (polyvinylidene fluoride) coated on either one or both sides with a conductive layer, were identified in Report 4 as the most suitable for the detection of plate waves in the aluminium sheets used to form the 'skin' of the Concept Demonstrator. Small-area sensors with an active area of 4.9 mm^2 (2.5 mm diameter circles) were affixed to the 'inside' of the aluminium sheets. This phase of the work looked at a number of different ways of attaching these sensors, with the aim of establishing a reliable and reproducible sensor suitable for 'mass production', as nearly 800 sensors were to form the sensor network.

Report 4 described sensors that consisted of a 10 mm diameter piece of PVDF film, 50 μm in thickness, with a silk-screen printed layer of silver-based (conductive) ink applied to one of the surfaces over circular area of about 2.5 mm in diameter. When glued to the aluminium sheet with a thin layer of epoxy, the aluminium sheet formed the ground electrode of the sensor, with the silver ink the active electrode. The voltage measured from the latter electrode provided a measure of the charge generated by the film during the passage of elastic waves past the sensor and hence through the aluminium sheet. During this stage of development, some abrasion of the silver ink electrodes occurred due to the spring-loaded pins that connect the sensors to the signal processing board.

A more convenient method of production of these sensors was thought to be to acquire commercially available gold-coated PVDF film (110 μm -thick film from Measurement Specialties Inc., MSI). The thicker film gave a higher voltage output from a given amplitude elastic wave, and the film could be punched out, using an in-house manufactured punch and die, to produce small discs 2.5 mm in diameter. These sheets were provided by MSI with gold coating on both sides. Again, the sensors were bonded to the aluminium sheets using a hard-setting epoxy (Epotek). These sensors formed the basis of a number of tests that sought to optimize the type of glue used to bond them to the aluminium, to investigate various protective layers applied to their active electrode to prevent wear by the spring-loaded pins, and to look at various ways of making electrical connections to them.

It has been found that both the silver-ink and the gold-coated PVDF sensors performed satisfactorily, although both coatings are subject to wear, especially the latter. This problem is

addressed below. The Epotek epoxy adhesive required both vacuum degassing to remove significant bubbles in the mixture (that form during mixing), and an extended curing time of over 24 hours. An alternative product requiring no degassing and which cured rapidly was Araldite 2015. The slightly more viscous mixture of this product was also easier to work with in small quantities and allowed a very thin layer to be applied.

The problem of wear on the electrodes of the sensors has been minimized by a new design of the mechanical structure to hold the electronics boards against the sensors (described below). However, under conditions where the Concept Demonstrator may be subjected to constant vibration, movement, or repetitive attachment and detachment of the panels, two possible solutions have been identified. The first is to insert a thin layer of anisotropically conductive silicone (Fujipoly WSC-035) between the spring-loaded pins and the backs of the sensors. This is easily applied (it may be held in place by a very small patch of double-sided tape or something similar) and does not modify or attenuate the signals from the sensors in any measureable way. The second is to silk-screen print a thin layer of Estapol (a polyurethane-based lacquer) to the back of the sensor. While this electrically insulates the sensor and adds a small series capacitance to the circuit, capacitive coupling yields a voltage signal essentially identical in form but reduced in amplitude (by about a factor of four) to that of the original signal. The polyurethane layer is highly resistant to any abrasion by the contacting pins. It has not been found necessary to implement either of these solutions at this stage. Any movement of the electronics boards (and hence the contacting pins), which may cause wearing of the sensor electrodes due to vibration during transport of the Concept Demonstrator, may be avoided by simply breaking the contact between the pins and the aluminium sheets by translating the boards using the structural mounting devices (height adjustment screws) described below.

3.3 Aluminium panels

The size of the aluminium panels that form the skin of the Concept Demonstrator is 220 mm x 200 mm. This represents an increase in size over the first set of panels used (200 mm x 200 mm), and was done for several reasons. Firstly, the height of the combined NAL and DAL boards was such that boards on panels adjacent to the internal corners of the hexagonal prism that forms the Concept Demonstrator would have made contact. Further, it was deemed desirable to have both non-square panels and sensors distributed over them not in a regular array. The arrangement of sensors and other features of the aluminium panels are shown in Figure 3.

The sensors were applied using Aradite 2015 as the adhesive, and were positioned using a purpose-built jig which held both the panel and a mask plate (with holes at the positions of the sensors) using two locating pins through the mounting holes of the aluminium sheet. Small divots were machined in the panel, two per sensor cell, to locate two of the earthing pins on the DAL board in order to ensure the sensor pins on the board were aligned to the sensors.

3.4 Mechanical fixtures

The first prototypes of the signal processing and networking electronics described in Report 4 have been miniaturized and consolidated onto two boards, each 80 mm square. These boards are referred to as the Data Acquisition Layer (DAL) board and the Network Application Layer (NAL) board. These two functional layers are also arranged in physical layers in the hardware




Figure 3: Layout of transducers, attachment holes and alignment divots on the 1-mm thick aluminium panels.

(Figure 2): the DAL board has spring-loaded pins that make contact with the sensors on the aluminium plate, and the NAL board is connected to the DAL board via a 40-pin high-density transition connector. These are described in detail in Section 4. Photographs from both sides of this composite module are shown in Figure 4. The spring-loaded gold pins can be seen in Figure 4(a): four of them contact the sensors, and the others contact the aluminium and are used for earth connection, shielding and, for two of them, board location. They can be seen more clearly in Figure 4(e), which shows the two boards mounted together.


The DAL board includes the electronics for four receivers and a single transmitter, although the board may be readily modified to be capable of transmitting on the receiving channels. Each of the active pins (either receiving or transmitting) has two earth pins nearby to reduce earth-loop interference. Therefore each sensor ‘cell’ has 15 spring-loaded pins that make contact with the aluminium panel, and with four such cells per panel a total of 60 pins are squeezed against the 1 mm-thick aluminium panel. The method of mounting these boards, with individual control of the force applied by each cell and without significant bending of the aluminium panel, is described in this section.

It was decided that the electronics boards of each cell should be mounted directly on the aluminium skin panel, and not from the frame of the Concept Demonstrator. Only the panel is mounted directly on the CD frame. The reason for this is to allow the cell hardware to be mounted, aligned, inspected and tested on the bench, before it is mounted in the frame. Similarly, if there is a malfunction, it can be removed from the frame for more convenient diagnosis of the problem.

To each cell (a DAL board attached to an NAL board) a small cross-shaped plastic (Ertalyte) fixture was affixed via four attachment points at the corners of the boards (see Figure 5). This cross had a central threaded spigot that in turn was attached via an adjusting nut and a lock nut to a larger cross made of polyurethane. (These were colloquially known as the small spider and the large spider, albeit spiders with the appearance of four legs!) The large spiders were attached to the aluminium sheet using five brass stand-offs, one at each end of a leg of the spider, and one at the centre. These roughly corresponded to the four corners and the centre of the aluminium sheet. Without the centre stand-off the aluminium sheet was found to bend excessively, requiring unnecessarily high forces to be applied to each individual cell to ensure that the pins near the centre of the sheet made contact with the corresponding sensor. On each brass stand-off a cam locking tab mechanism was attached that allowed the whole assembly to be quickly affixed to the frame by turning four locking tabs that mate with four striker plates on the frame (see Figure 6).

This light yet strong and rigid ‘double spider’ structure allowed each cell to be raised and lowered with respect to the aluminium sheet as necessary. The cells could be locked in an ‘up’ position (with the cell not in contact with the sensors) for transport or to avoid unnecessary wear on the sensors, or compressed against the aluminium sheet and the sensors with minimal force (depending on the tension in each group of spring-loaded pins in each cell) and locked in position. The location of the pins with respect to the sensors was facilitated by machining two small divots in the aluminium sheet which located two diagonally opposite earthing pins, therefore accurately aligning the active sensor pins over the sensors themselves.

(a)




(b)



Figure 4: (a) Lower side of DAL board, which faces the aluminium sheet, and (b) upper side of DAL board. The pins that connect to the sensors are the central pins in each group of three: those at the corners of the board connect to the PVDF sensors, while the group of pins at the center of the board is for connection to a piezoelectric transmitter. The DSP and the four sensor amplifiers (2 each on the left and right-hand sides) can be seen in (a). The communication link to the NAL board is carried by the large (40-pin) connector at lower center in (b).

(c)



(d)




Figure 4: (c) Lower side of NAL board and (d) upper side of NAL board. The communication link to the DAL board is the large (40-pin) connector at the top of the image in (c). The four serial ports that connect to the neighbouring cells can be seen in (d). The two chips to the right of centre in (d) are the FPGA (with the green spot) and the DSP (see Section 4 for information on these components). Both are ball grid arrays.




Figure 4: (e) DAL (lower) and NAL (upper) boards connected together.






Figure 5: The DAL and NAL boards connected, with the small 'spider' attached, showing the height adjustment and locking nuts on the central mounting screw.



(a)



(b)



(c)

Figure 6: (a) Assembled panel showing the large 'spider' attached to the aluminium panel and holding the four smaller 'spiders' and their DAL and NAL boards in place. (b) The inside 'skin' of the Concept Demonstrator showing the array of panels. (c) CAD drawing of the assembled panels showing the main mechanical features of the support structure for the electronics

3.5 The Concept Demonstrator

The hexagonal prism that forms the Concept Demonstrator is mounted with its six-fold axis vertical, and measures 838 mm high (including the base plate), 883 mm across the vertices and 764 mm from the one of the hexagonal faces to its opposite face. Each face is approximately 440 mm x 800 mm. As described in earlier reports, its frame is made of extruded aluminium (Maytec Profile 20 mm x 20mm, 4H, LP), which basically consists of a square profile section (20 mm) with a slot along each face. Into this slot we have laid sponge rubber round cords, so that around three edges of each aluminium panel there is a compliant seal that reduces vibrations being transferred between panels and prevents panels rattling against the frame.

Panels with the four cells attached are inserted into the frame from the outside, and locked into position as described in the previous section. Each cell is connected to its nearest neighbour via four 20-pin IDC micro-ribbon cables. The entire Concept Demonstrator is mounted on a manually controlled x-y table, which is in turn mounted to a hydraulic ram allowing z motion. The Demonstrator may be locked in position to prevent rotation and translation, or may be moved about 210 mm in the x-direction, 100 mm in the y-direction (the direction of the laser beam), and 310 mm in the z-direction. It is also free to rotate about its hexagonal axis. Rapid firing of the laser combined with motion in any of these directions may be used in the future to demonstrate the systems' response to multiple hits. Photographs of the Demonstrator during assembly are shown in Figure 7.

3.6 Determination of impact location

At the present stage of the project the determination of the location of an impact has been calculated using triangulation based on measured arrival times of the lowest order extensional wave and the known group velocity of these waves at a particular frequency. At present, after the time-delay calculations, a look-up table of values is used to find the impact location. This method is fast and uses very little memory space on the DSP. The method of determining the location, and the algorithm that generates the look-up table, are described below. An analysis of impact location by triangulation, and the uncertainties involved, is included as an Appendix to this Report.

Signals from the four sensors are detected by the DAL electronics, narrow-band filtered, amplified and digitized. For each of these four channels, the time at which the first signal exceeds a preset threshold level is found. The sampling rate of the digitizer is 3.125 MS/s, giving a time step interval of 320 ns. The earliest arrival time is then subtracted from the other three times to give three time delays, which are used to estimate the location of the impact relative to the centre of the square formed by the four sensors used. The data processing results in the measurement of the arrival time of the lowest order extensional plate wave (S_0) centred on a frequency of about 1.5 MHz, and therefore with a group velocity in the aluminium plate of about 5300 m/s.

If standard triangulation techniques were applied to these three time delays, the equations for three hyperbolae would need to be solved simultaneously. Further, as the digitization rate is limited to around 3 MHz in the present hardware, the time of arrival of the extensional waves may not be determined accurately enough for solutions to these equations to exist. Searching for near solutions can be done, but this would take a significant amount of processor time and





Figure 7: Three stages of assembly of the Concept Demonstrator.

memory in the present configuration. In the present version of the Demonstrator a faster way has been to use a look-up table that has been generated as follows.

As each cell has four sensors arranged at the corners of a 60 mm square, the look-up table will be the same for each of the eight octants whose origin is at the centre of the square. For the present size of aluminium panel the largest area that needs to be covered by the look-up table is a truncated, isosceles (at 45°) right triangle (shown in Figure 8(a)), 165 mm long and 143 mm in the truncated direction. This geometry will cover the ‘worst case’ of an impact in a corner of the panel diagonally opposite that of the cell containing the four sensors being used. This area was then divided into 1-mm squares. For each square the time delays may be calculated. These delays are used to form a six-figure ‘index’ that is associated with each position. This index number is formed from the three delay times from each pair of sensors, the first two digits being the shortest delay time (in number of time step intervals of 320 ns), the second two-digit number is the next longest delay time, and the final two-digit number the longest delay time. While the total number of points (on a 1-mm grid) inside this truncated triangular area is about 13600, the total number of unique points (points that have different index numbers) is only 1071 due to the finite time-step interval. The mean position of all the 1-mm cells that have the same index number is stored in the look-up table with that index number. The look-up table is stored in the Flash memory on the DSP.

When an impact is registered, the delays are first ordered to quickly determine which octant of the Cartesian plane contained the impact site. Depending on the order of the impacts (the time of arrival at each of the sensors) one of eight co-ordinate transformation matrices is used to translate the look-up table result to the correct orientation. The index number of the impact is then found in the look-up table and the average position associated with that index number multiplied by the appropriate transformation matrix is given as the location of the impact. In this way any of the four sets of four sensors on a panel can detect the location of an impact. A primitive form of error estimation and impact severity has also been implemented.

Figure 8(b) illustrates this process. In the picture each pixel represents one of the 1-mm square areas. Each green point is the average position of all of the 1-mm squares that have the same index number. The other colours are a representation of the total x- and y-coordinate errors, with red shading for y-errors and blue for x-errors. The two yellow rhombs drawn are to aid the eye in visualizing the area covered by the index number. The further from the sensors the sparser are the index points and the larger the positioning uncertainties. This is, of course, due to the fact that the time differences far from the sensors make it difficult to locate the impact precisely, and this is worst outside the square formed by the sensors, and on a diagonal of that square. The error boundaries have been represented by rhombs, as if the calculation were done analytically solving for intersecting hyperbolae, the error boundaries would approximate these shapes. The analytic results given in Appendix A show how the location uncertainties increase dramatically (or, equivalently, the geometric dilution of precision increases) when the impact is outside the region bounded by the sensors.

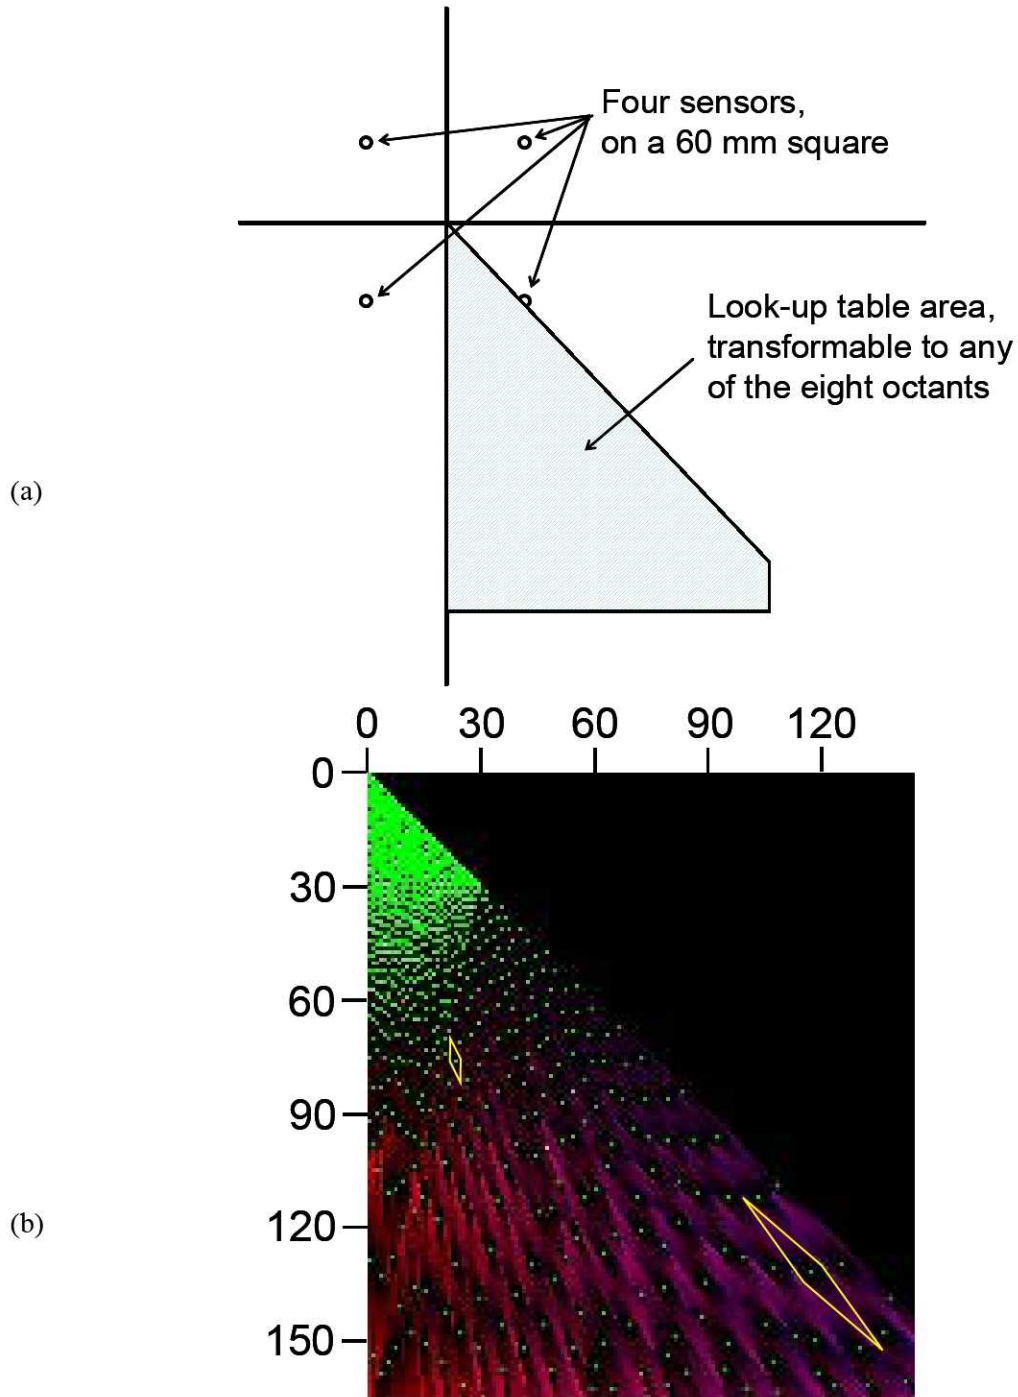


Figure 8: (a) Position of sensors and the area of the panel covered by the look-up table. The table is transformed to the other eight octants to cover the full panel. (b) Graphical representation of look-up table points and estimates of the positional uncertainties. The top left-hand corner represents the centre of the group of four sensors, and the axes' labels are millimeters from the centre of the square formed by the sensors. The green dots are the points in the look-up table and the yellow rhombs represent the boundaries of areas covered by these points, and therefore give a representation of the uncertainties in the estimate of the location of an impact.

4. Hardware and Software Implementation of the Physical Cells

In the System Overview (Section 2) it was explained that cells consist of sensors, a Data Acquisition Sub-module (DAS) and a Network Application Sub-module (NAS), and these are described in this Section. Point-to-point communications links are provided in the Network Application Layer by serial communication links between network application sub-modules, and in this Section the communications protocols for these physical links will also be explained.

4.1 Data Acquisition Sub-Module

The Data Acquisition Sub-module (DAS) is connected to the four sensors in the cell, and the purpose of the sub-module is to detect impacts from the signals received from each sensor. More specifically, the DAS is designed to accept signals from four PVDF sensors, condition them, sample the signals into a digital signal processor (DSP) system, and compute estimates of impact position and signal level. This information is then communicated to the attached Network Application Processor.

The prototype of the DAS was implemented in three sub-layers to speed up development. The layers performed:

- analog amplification and filtering, which was performed immediately adjacent to the sensors,
- anti-alias filtering and signal offsetting for the DSP sampling system, and
- DSP-based sampling and processing, implemented on a commercial DSP Starter Kit.

In the current version of the DAS there have been two major changes. Firstly, all three sub-layers have been integrated into a single circuit board, placed between the sensors and the Network Application Sub-module (NAS). Secondly, the system has been enhanced by the provision of an output channel capable of generating a stimulus into the skin for diagnostic purposes. The output channel requires an additional piezoelectric transducer on the skin, which is only implemented on some of the skin panels. These major changes and some minor changes will be described below.

DAS System Integration

The integrated version of the DAS board is equipped with spring-loaded pins to contact the sensors on the aluminium skin panel, signal amplifiers and filters to condition the analog signals, and the DSP and control functions, which enable impact position and signal amplitude determination. The whole system is implemented on an 80 mm x 80 mm, 6-layer printed circuit board as shown above in Figures 4(a) and (b). The basic design places the DSP, control circuits and power supplies in the central area of the board, and the analog processing around the edges.

The analog receivers consist of a second order high-pass filter at about 350 kHz, a fifth order Butterworth low-pass anti-alias filter at about 1.55 MHz, and signal amplification and offsetting circuits. The amplifier outputs are passed to the DSP's analog inputs. The system has also been equipped with external hardware comparators, which have programmable thresholds that will be subsequently tested as potential digital input channels to aid with positioning calculations.

The original DAS design was intended to use the Texas Instruments TMS320F2810 processor. Production problems at TI resulted in delivery delays and, subsequently, both the substitution of the TMS320F2812 processor and the substantial scaling back from a fully populated Concept Demonstrator with 4 x 192 sensors to a smaller array of sensors. Failure of PCB assembly equipment at the selected supplier has further curtailed the size of the sensor array that can be immediately demonstrated. Consideration will be given to the optimum number of sensors to demonstrate CD functionality in the immediate future. Ongoing developments will probably consider other sensing options and modalities rather than just an increase in DAS numbers.

The connections for almost all of the digital signals have been routed through a CPLD¹, which allows almost complete redefinition of the DSP's digital connections to the on-board hardware and the connection to the NAS. The CPLD can be dynamically configured from the NAS. The DSP can have new software downloaded from the NAS as well. In this way both the software and much of the hardware for the DAS can be modified without removal from the system. In addition the CPLD allows for the definition of extra logical functions, such as access from the NAS to the DAS's timing signals, and even direct calculation of impact time of arrival delays for the position determination.

DAS Acoustic Signal Generation

The use of a large number of sensors creates testing and calibration problems: there are significant variations in the sensitivities of the sensors in some cells, largely as a result of bonding variations, and it is possible that sensitivities may change with age. The DAS has therefore been provided with the capability of generating an acoustic signal from a centrally located piezoelectric disc (at this stage, PZT), to provide a means of monitoring and calibrating the sensors. This facility may also be useful for active damage evaluation.

As can be seen in the board designs above (Figures 4(a), (b)), there is an additional pin set mounted in the centre of the board. This pin set is driven from a buffer amplifier connected to the CPLD. The amplifier allows a 10 Vpp binary waveform to be applied to a piezoelectric transducer in the centre of the 4 sensors. The applied signal can be either a fundamental square-wave or a more complex pulse-width modulated binary signal. The signal can be sourced from the CPLD itself, the DSP, or even the NAS via the NAS-to-DAS interface. A subset of the sensor panels has been modified to support the PZT transducers. This function will not be detailed in this report, but will form part of the experimentation for the next stage of the Concept Demonstrator development.

4.2 Network Application Sub-Module

The Network Application Sub-module (NAS) receives sensor and status information from the attached DAS, is the platform on which the agent processes run, and provides the network that allows the agents to communicate with each other.

The major components of the NAS are as follows.

- Texas Instruments TMS320VC5509, which is a 288 MIPS fixed-point digital signal processor (DSP).
- 8 Mbytes of volatile memory (SDRAM).

¹ Complex Programmable Logic Device – a Cool Runner II, XC2C256 from Xilinx.

- 2 Mbytes of non-volatile memory (Flash).
- High-speed synchronous serial communications link to the DAS (using the McBSP² built into the DSP). The communications software is implemented using a simple fixed-length packet for the transfer of data.
- Xilinx Spartan II, XC2S400E Field Programmable Gate Array (FPGA), which controls the four asynchronous serial communications links to the four neighbouring cells in the network array.

The board is shown above in Figures 4(c) and 4(d).

There have been important changes in the implementation and features of the NAS since the Phase 1 report (CTIP, 2003b). The major architectural change has been the replacement of the original CPLD and UART³ communications interface, with an FPGA-based interface. This allows:

- faster communications between the NAS boards,
- flexible definition of the board's functions, and
- the transfer of some aspects of the data communications protocols from the DSP software to external hardware. (This is expected to allow the development of much more sophisticated communications routing protocols.)

Another enhancement is the development of interface current monitoring. As the present power supply arrangements obtain power from the edge of the array, a single major cell failure could overload the whole system. The new NAS monitors currents on each of its four neighbour interfaces. The system uses the DSP's analog input pins to measure each of the four currents. Full scale is $\pm 3A$, measured in approximately 5 mA steps. This facility allows the system to detect problems with power supply loading.

The FPGA on the NAS loads its configuration file from the DSP unit. This means that the hardware function can be modified and a new configuration loaded at any time. This is a great advantage for the testing of new hardware protocols.

Over 192 of these NAS modules have been built and tested, allowing a fully populated NAS array to be demonstrated (see Figures 7 and 9).

4.3 System Integration

The NAS and DAS are connected to form a single cell of the demonstrator. Four cells are mounted together to form a panel of the system as shown in Section 3 in Figure 6(a). The four-cell panels are mounted as shown in Figure 9 (see also Figures 6(b) and 7 in Section 3) to form the outer skin of the Concept Demonstrator. New connectors have been used to ensure more reliable connections. The basic power supply system can be seen at the bottom of the Concept Demonstrator. The power distribution and communications links can also be seen at the top and bottom in Figure 9.

4.4 Software and Configuration Updates

The cell architecture, including both the NAS and DAS, has been designed to allow easy updating of software. The normal procedure is that each NAS DSP is initially loaded with

² Multi-channel Buffered Serial Port

³ Universal Asynchronous Receiver Transmitter.

code that allows it to load code upgrades. After the initial NAS DSP code is installed, subsequent upgrades occur as follows.

- The network can be used to deliver code upgrades to the NAS DSPs.
- Once programmed, the NAS DSPs can load new configuration data into their FPGAs.
- The NAS DSPs can then reconfigure the DAS CPLDs via the FPGA circuits.
- The NAS DSPs can then force a reboot and new code installation on the DAS DSPs.

In this way, after code is distributed to the NAS DSPs, all other upgrading takes place in parallel. The only process that requires each processor to be dealt with individually is the initial NAS DSP code installation. All NAS DSPs have been through this process.




Figure 9: The Concept Demonstrator with four of the six faces populated with NAS units.

DAS Software and Programming Procedures

At the core of the digital electronics on the Data Acquisition Sub-Module (DAS) is a Texas Instruments TMS320F2812 Digital Signal Processor (DSP). This DSP continuously samples four analog channels at 3.125 MSPS (Million Samples Per Second) through its built-in Analog to Digital Converter (ADC). These digital values are stored in a circular buffer containing 200 samples, or 64 μ s of data, from each channel. Every fourth sample from each channel is tested to check if it has deviated by more than 90mV from the channel's average value. Once a channel has exceeded this selected threshold, a further 184 samples are taken on each channel, an impact is flagged, and the buffers are processed. Each channel is processed to give the relative time that the threshold was passed. The maximum deviation from the average in a 10-sample window around the crossing of this threshold is also recorded. This provides an initial estimate of signal amplitude. The relative time delays are then used to index a look-up table, which provides an estimate of the location of the impact that caused the signals (described in Section 3). This information is then sent to the NAS using the McBSP communications.

The McBSP communications protocols have not changed in the present implementation but the communications capability has been enhanced. The Phase 1 demonstrator (CTIP, 2003b) allowed communication only from the DAS to the NAS at a burst rate of 500 kbits/s, while the current software has duplex communications between the DAS and the NAS. The data rate from the DAS to the NAS is 6.25 Mbits/s and 6 Mbits/s from the NAS to the DAS.

The Xilinx Complex Programmable Logic Device (CPLD) is used to make all digital connections on the DAS. This includes the communications between the DSP and the NAS. The channels that carry the McBSP communications from the NAS to the DAS can also be used to communicate to the DSP's Serial Peripheral Interface (SPI), which is used when reprogramming the DSP. The CPLD is also used to enable the four analog amplifiers, which filter and amplify the signals from the sensors to the DSP ADC.

Both the CPLD and the DSP are re-programmable via the NAS. Reprogramming the Flash memory built into the DSP is a two-stage process. The boot ROM on the DSP contains a number of selectable boot modes including a boot from SPI mode. The NAS first sets the signals to enable the SPI communications, then resets the DSP in the boot from SPI mode. The DSP on the NAS contains code so that it appears as a boot PROM with an SPI-Slave interface. The boot from SPI code contained in the ROM is unable to load data to the Flash memory, so it is sent code that will run from RAM. Once all the RAM code is received it is executed. This program was deliberately written to continue the SPI communication and includes code to write to the Flash memory. Once the Flash has been erased, the SPI communication continues and the actual DAS program is sent along with the triangulation look-up table to be stored in Flash. When the DAS signals the SPI communication has finished, the NAS configures the DAS to reboot from Flash and enables McBSP communications. This completes the DSP reprogramming process.

The CPLD can be reprogrammed via an IEEE 1149.1 JTAG interface. The NAS DSP is used to emulate a JTAG interface. The source code required to implement this has been specially modified from a program written to run on a microprocessor. The programming file is in binary XSVF (Xilinx Serial Vector Format) and is generated from the design program.

4.5 Network Communications and Protocols

The network implemented in the Network Application Layer (NAL) is used for a number of purposes, some of which are only for use in the test-bed and would not be present in a deployed IVHM system. The main communication categories are:

- communication between agents (or other processing) running on separate NAL modules to implement the distributed system intelligence and initiate actions,
- distribution of software to NAS and DAS modules, and
- transmission of diagnostic information for visualizing the status of the test-bed and studying its behavior.

In this section an update on the communications protocols used in the sensor network (i.e. between Network Application Sub-modules) is presented. The communication stack is layered, with a physical layer, a data link layer and a network layer (the lowest three layers of the OSI Reference Model (Tanenbaum, 2003, p. 37)) being implemented. The higher layers are at present not separate communications protocols but are subsumed into the application code. Although the lowest two layers have not changed since the previous Report (CTIP, 2003b), further details are provided which are required to understand the significant changes in the network layer protocols.

Physical Layer

The four communications ports in each Network Application Sub-module contain eight wires, four in each direction. These are connected directly to a field programmable gate array (FPGA), which allows the hardware interface to be reconfigured. The interfaces have been configured as high-speed asynchronous serial interfaces, similar to RS-232, but with differences noted below. The signals are directly connected between FPGAs on boards without voltage translation, so the signals are 0V and 3.3V rather than the RS-232 bipolar voltages.

Of the eight available wires in the physical interface, only six in three pairs of signals are used.

- TxD (Transmit Data) and RxD (Receive Data) – the signal follows the RS-232 standard, with 8 data bits, 1 stop bit and no parity bit.
- RTS (Request to Send) and CTS (Clear to Send) – The RTS output signals the state of the receive buffer, and the CTS input stops transmissions from TxD when the buffer is almost full.
- DTR (Data Terminal Ready) and DSR (Data Set Ready) – The DTR output signals that the communications link is active, and the loss of the received signal on DSR indicates a link failure.

The remaining two wires will be used for accurate time synchronization between cells. The transmission speeds on the serial interface range from 9600 bits/s (base rate) to 1.5 Mbits/s, and each communications port has a 256 character FIFO (first-in-first-out) buffer on both the input and output. In all, this employs only 12% of the FPGA logic, so it will be possible in the future to move some of the data link and even network layer protocol processing from the DSP to the FPGA if necessary. It would also be possible to change the communications from asynchronous serial to synchronous serial at up to 48 Mbits/s without any hardware changes should the need arise (although the DSP would not be capable of processing the packets if this rate were sustained on even one of the ports).

Data Link Layer

The data framing used for the data link layer of the point-to-point communication links is a simplification of the PPP format (Simpson, 1994a,b), and is shown in Figure 10. The protocols will be discussed in the following sub-section, the payload can be up to 1500 bytes, and the checksum is the same as the 16-bit checksum used by Simpson (1994a). Byte stuffing is used to prevent the flag character appearing in the payload.




Figure 10: Serial data link frame format.

In addition to the physical cells, there are also simulated cells, and as there are many simulated cells per workstation it is not practical to have enough serial ports to allow these cells to be interconnected via serial links. Thus for communications links between simulated cells, i.e. between cells simulated on the same and other workstations, the data frame is transported using UDP (Tanenbaum, 2003, pp. 524–532), with the UDP payload consisting of the protocol, payload and checksum shown in Figure 10. This difference is hidden from the higher levels of

the communications stack. Links between workstations and physical cells use high-speed (up to 1.5 Mbits/s) serial links, the same as between physical cells.

When a serial communications link is first formed, negotiation is required to determine the serial data rate that will be used, which is initially performed at the base rate (9600 bits/s). When the link is established there are also regular packets to verify that the link is open and to monitor link delays. This is the function of the Link Control Protocol (LCP). There are only four LCP messages as follows.

- *CapReq* – Capability Request, this contains the capabilities (e.g. serial speeds) of the sender, and is used to negotiate the communications parameters.
- *CapAck* – Capability Acknowledgement, this is a reply to a *CapReq* message, and contains the selected communications parameters (e.g. the highest common serial speed).
- *EchoReq* – Echo Request, this is used to confirm that the link is open, and contains the sending time (millisecond resolution). The receiver can use this for coarse time synchronization.
- *EchoAck* – Echo Acknowledgement, this is sent in reply to an *EchoReq* message and contains the sequence number and send time in that message. Besides confirming link operation, this can be used to estimate link delays.

The MAC address is also included in each message. Agreement is reached on the communication parameters when a *CapAck* message has been both sent and received with commonly agreed communications parameters. Then, after a delay, the serial link speed is changed and the link is tested. When an *EchoReq* has been both sent and received (within a timeout period) the link is open for use by the higher layers. This exchange is illustrated in Figure 11. The protocol also provides a mechanism for renegotiating a lower link speed if the link test fails. The *EchoReq* message is periodically sent while the link is open, and if no reply is received after a few messages have been sent the link is closed to higher-level protocols until a connection is re-established.




Figure 11: Link Control Protocol exchange to negotiate link speed and test link.

Network Layer Protocols

The data link layer supports multiple network layer protocols (up to 255), and a range of protocols will be required to support the different types of interactions between nodes. Consideration of the type of communications interactions that will occur in the Concept Demonstrator shows that they can be divided into the following types.

1. **Concept Demonstrator Configuration:** Send packets from test-bed visualization workstation to some or all cells in the Concept Demonstrator. This is typically used to configure the network or update cell code.
2. **Network Visualization:** Send packets from a cell to the test-bed visualization workstation. This is typically used to report data for visualization on the workstation.
3. **Agent Communication:** Exchange data between agents running on cells. This will only occur between cells that are either directly connected or in close proximity.
4. **Network Routing:** Communication between a cell and a sensor network service, which could be at any distance. The service could be simple, such as a time service that is a standard cell elected to the position, or tied to a physical entity such as the vehicle propulsion system. This would also include the home service of a mobile inspection agent.

The first two types of interaction are specific to the test-bed and would not occur in an operational sensor network. For robustness the visualization workstation is physically connected to several cells well-spaced around the periphery of the Concept Demonstrator, and this will have some bearing on the design of the protocols.

Important characteristics required of all types of communication are that they be robust and scalable, and that there is a graceful degradation in the network performance with substantial damage to the sensor network. An important reason for using multi-agent algorithms for distributed processing is for scalability, as these algorithms only have local interaction, which remains constant as the network size increases. This will be the predominant form of communication in an operational network, with communication with services containing highly summarized data. Interactions of type 2 will not scale well, as larger networks result in more cells reporting their status, which in turn increases the data going to a specific node (the test-bed visualization workstation). This is not a major problem, as this protocol is not used in an operational sensor network, so the reported information needs to be limited so as to operate in the test-bed. The implemented protocols for each type of interaction are described in the following paragraphs.

Concept Demonstrator Configuration

Packets sent from the visualization workstation are typically sent to all cells in the network, although in some (the minority of) cases they are directed to specific cells. The simplest and most robust method to deliver the payload to all cells in the network is to use flooding (Tanenbaum, 2003, p. 355), and the selected packet format is shown in Figure 12, with the protocol being called the Master Flood Protocol (MFP). The source of the packets is not identified as there is only permitted to be a single visualization workstation in the sensor network test-bed. Packets are sent from the visualization workstation with the same sequence number, and when received by a cell, packets are either discarded, if they have been previously received, or they are sent out to all ports other than the one on which they were received, in addition to being locally processed. The *Hops to Live* parameter (also called *Time to Live*) prevents packets circulating indefinitely in the network.




Figure 12: Packet format for Master Flood Protocol for packets sent from test-bed visualization workstation.

Network Visualization

Packets from individual cells destined for the visualization workstation are routed through the sensor network. The routing protocol, called the Master Route Protocol (MRP), can be relatively simple as there is only a single destination for all such packets, but there are multiple destination ports on the workstation. To allow control over the robustness of packet delivery versus the amount of network bandwidth consumed, packets can be routed to the N closest (counting number of hops) ports on the workstation, where N can range from one to the number of workstation ports connected to the sensor network. As the cells to which the workstation connects are well distributed around the edges of the sensor network, the packets take quite different routes through the network. It is therefore expected that the system should be robust to the loss of packets due to damage or congestion.

The routing table in each cell specifies the distance (number of hops) to each port of the visualization workstation, and the port on the local cell through which the packet should be sent. The packet format is shown in Figure 13. The source MAC address is the unique 64-bit identifier in each cell. The routing table is constructed by monitoring packets from the test-bed visualization computer sent using the Master Flood Protocol.

When a packet is sent through a simulated cell, it is sent directly (i.e. in one hop) to the visualization workstation using UDP, rather than being routed via multiple hops through a chain of simulated cells. To allow for this when building routing tables, all simulated cells have a distance of one hop from the visualization workstation. This is done to reduce the computational workload of simulated cells.




Figure 13: Packet format for Master Route Protocol for packets sent to test-bed visualization computer.

Agent Communication

There are two protocols that are used for data exchange between agents. For all of the multi-agent algorithms that we are currently using the communication is only between directly connected cells. The protocol that was created for this is called the Simple Agent Protocol (SAP), and is mainly a wrapper on the data layer frames. This is used so that a consistent software interface is presented by all network protocols to the higher layers of the communications stack. Another function of the SAP is that random errors can be induced to simulate damage resulting from impacts. This is used for evaluating the robustness of the multi-agent algorithms, although clearly such intentional data corruption would not occur in an operational network.

The other situation for agent interaction is communication with other agents within a fixed distance, typically only two or three hops. For this purpose a flood protocol, which is called

the Concept Demonstrator Flood Protocol (CDFP), may be used. This is a separate protocol from the Master Flood Protocol, not only to keep separate the operational and test-bed protocols, but also because the CDFP needs to allow for multiple sources of packets, rather than just one. The packet format is shown in Figure 14, with the difference from the MFP being the inclusion of the source MAC address. Processing the CDFP packets is similar to the MFP, but the MAC address as well as Sequence Number of recent packets must be used to determine which packets have been previously received, and hence should be rejected. As the communication range is fixed (set by a multi-agent algorithm), irrespective of the size of the sensor network, the communication load on each cell does not depend on the size of the network, so both protocols will scale to any size sensor network.




Figure 14: Packet format for Concept Demonstrator Flood Protocol.

Network Routing

The final type of interaction requires communication between cells at any distance, but not between arbitrary cells. In all cases of long distance communication at least one of the cells will have some type of special function (it may also be an ordinary sensing node as well). The approach to the routing for this sensor network is to use hop-by-hop routing, as is commonly used in *ad hoc* networking algorithms such as AODV (Ad hoc On-demand Distance Vector) routing protocols. This is simpler than source routing where the node must know the entire route to the destination. Furthermore, a hybrid between a table-driven proactive approach (in which the routing table is continuously maintained) and an on-demand reactive approach (in which routes are found only when needed) will be used. (See Ilyas, 2003, Ch. 21, pp. 2-5 for further details on *ad hoc* routing algorithms).

The implementation divides the operation into two protocols:

- i) a data forwarding and error control protocol, similar to IP and ICMP (Tanenbaum, 2003, p. 449) in functionality, and
- ii) a route control protocol responsible for route discovery and maintenance.

The data forwarding protocol forwards a packet using the entry in the routing table, if it is valid; otherwise it requests a route using the route control protocol. If the packet cannot be forwarded due to an error then it is sent back towards the source of the packet, with an error bit set, until an alternative valid route is found, or the source cell is reached and the packet is discarded.

The route control protocol is still in development. The initial implementation uses flooding from the services to build routing tables in each cell for hop-by-hop routing back to the service. This is a purely table-driven proactive approach, and is adequate for current needs, but this protocol will not scale well and is not particularly adaptive to network changes. An improved protocol currently being explored is to have dynamic clustering and election of cluster heads. When a cell does have a current and valid route to a source, it requests the route from its cluster head. If this is not available the cluster head finds a route dynamically using an approach similar to AODV routing. The eventual form of the route control protocol will evolve with the types of services that are offered in the sensor network.

5. System Operation Workstation and Visualization

A PC workstation is used to operate and monitor the test-bed. It is important to make clear that this workstation does not control the operation of the system when the agent software is running: in this case it acts simply as a display system that provides a visual output of the state of each cell.

The system has been designed so that this workstation can control the network functions by being attached to either one or more of the physical cells (using serial communication via a USB port) or one or more of the simulated cells (using UDP communication). This functionality is being expanded upon with further development of the test-bed, and current operations include the following.

- Distribution of code to the cells, for both the NAS and DAS processors.
- Initialization of the clock time on the cells.
- Dynamical determination and display of the network topology status of each link, where the addition or removal of cells or links are displayed in real-time.
- Dynamic display of the current flow to and from each cell, through each port, for diagnostic purposes.
- Real time display of impact locations.
- Display of cells suffering damage, whether due to impacts or other sources of failure.
- Display of agent processing state (described below).

5.1 Visualization software

Building the Network

Starting from a point where the workstation is connected to the network of cells, the concept demonstrator visualizer, *CDVis*, dynamically builds the network of simulated and physical cells by probing each cell to find the status of its communication links to its neighbours, and then probing each of its neighbouring cells. This recursive method allows a topographical interpretation of the entire network array to be built showing the connectivity of every cell to its neighbours (Figure 15). At this stage a two-dimensional image of the hexagonal prism array is shown: the left- and right-hand ends of the array image are connected together, as indicated by the green status of the communication links in the Figure.

Current Monitoring

At present, each cell is connected to each of its four immediate neighbours in a square array, and these connections carry both neighbourhood communications and power distribution. As each cell can draw in the vicinity of 200 mA when both NAL and DAL boards are operating, a significant amount of current (up to ~ 40 A) is drawn by the entire network. To monitor the current flow through the system, the visualizer has a function in which it periodically polls the cells for their current measurements on each port, using the current monitoring facility outlined in Section 4.2. Each cell then responds with its current measurements, after a small random delay to prevent buffer overflows that may result in lost communication packets. The visualizer can display a view of this information (Figure 16): it provides an intuitive representation of the current flows throughout the system. This view is dynamically updated such that the visual perception is that there is continuous flowing of the pixels, which allows

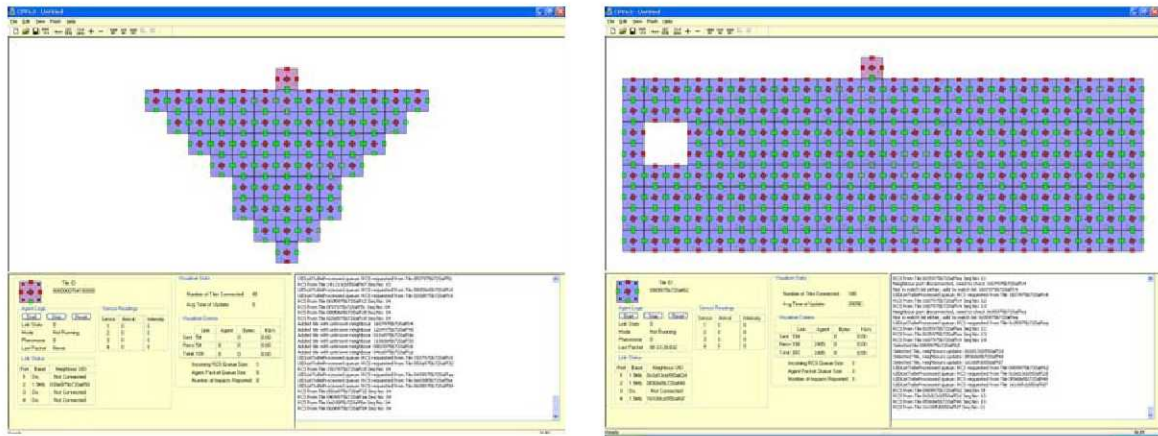


Figure 15: Network discovery by the visualizer of 188 physical cells, as shown on the visualizer screen. The visualizer workstation is depicted as the cell shown in pink. The screen on the left shows the network discovery process in operation, while that on the right shows the complete network. The four missing cells at the left hand end of the array (right image) is a ‘window’ left unfilled to allow physical observation of the inside of the demonstrator (see Figure 19 in Section 7).

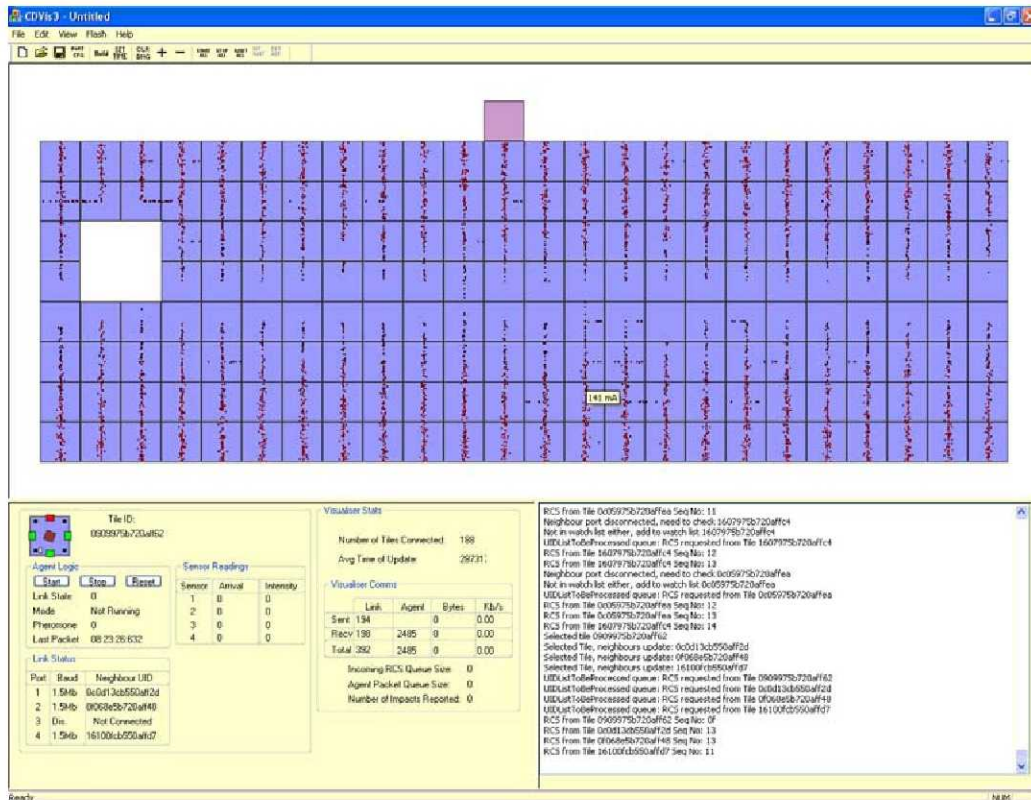


Figure 16: Current Flow View

for immediate discernment of the direction of current flow. In addition to current flow direction, the current value is represented by increasing redness in colour and thickness as the current increases. Notice that the current is naturally routing around the four missing cells on the left in Figure 16. The actual current value for any port can be displayed by simply placing the mouse cursor over the port (shown at the lower center of the figure).

Cell State Visualization

Each cell is capable of running multiple algorithms, and each algorithm incorporates the option of sending state information to the visualizer at any cycle. This is useful for providing a visual indication of how the multi-agent system is working, and particularly useful for debugging purposes. Traditional debugging of programs allows the programmer to trace through the code line by line. However, in an asynchronous multi-agent system, where the current state of one agent may depend on a number of other agents that may all be in different states, debugging is more complex. Tracing through the code of a cell may affect the cell's interactions with neighbouring agents, which in turn may lead to the cell following a different program path. By displaying visual representations of the states of the cells, an assessment of the operation of the entire system can be deduced quickly.

Impact locations, as detected by the cell sensor arrays, and their uncertainties are determined as outlined in Section 3.6 (see also Appendix A). These are shown on the impacted cell as a rectangular box, as shown in Figure 21, Section 7.

Impact Boundary Display

Report 4 (CTIP, 2003b) described the emergent formation of impact boundaries that delineate damaged regions caused by particle impacts, and this algorithm has been implemented in the Concept Demonstrator. Logic state information about the self-organizing boundary is sent to the visualizer from each cell whenever there is a change of state. This allows the observer to see the temporal formation of emergent boundaries, which result from purely local neighbourhood information exchange.

The impact boundary agent algorithm results in the formation of boundaries around critically damaged areas of the demonstrator. As it would be an expensive exercise to test this logic by critically damaging cells for each test, simulated damage that causes communication data corruption has been introduced. This will be described further in Section 7 below. (See also Section 4.5.)

Impact Network Display

An algorithm that produces a network connecting sub-critically damaged cells was also described in Report 4 (CTIP, 2003b). This algorithm simulates the behaviour of ants foraging for food, and can result in a system with nearly all agents in different states due to the deposition and evaporation of pheromone. Therefore, visualization information pertaining to this algorithm is sent periodically to provide information about pheromone levels and ant locations. Resulting views are shown in Figure 17. These impact networks may facilitate the fast inspection of impact areas, the identification of high densities of impacts, and the routing of repair resources.

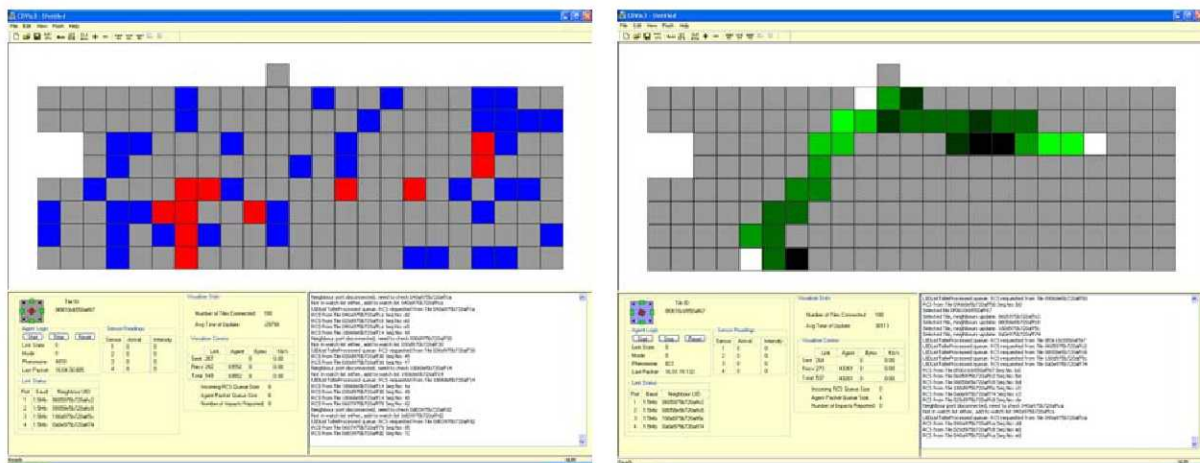


Figure 17: Visualization screen views for the impact networks algorithm. Ants view (left): Blue ants are exploring ants; red ants have discovered a food source and are returning home. Pheromone view (right) with non-critically impacted cells shown in white, and the pheromone levels in other cells shown in green (the greener the cell, the more pheromone present).

6. Intelligent Distributed Processing

The solution that has been adopted for the problems of handling the vast quantities of data and making the IVHM system robust, is to distribute the processing so that most computation takes place near the data source, and there is no single point or small number of points of failure resulting from the use of central computers. In other words, without centralized controllers, agents (cells) are expected to self-organize and survive on the basis of local, rather than global, information – no single agent has access to information about any others but its neighbours. This approach uses the idea of localized algorithms, in which simple local node behaviours achieve a desired global objective, while communicating only with nodes within an immediate neighbourhood.

Our proposed approach to self-organizing multi-agent networks is based on an iterative process including the following steps:

- a) forward simulation leading to emergent behaviour for a class of localized algorithms;
- b) quantitative measurement of spatio-temporal stability of the emerging patterns, using information-theoretic metrics for phase transitions; and
- c) evolutionary modelling of the algorithms, with the fitness functions corresponding to the metrics obtained at step b).

The last Report (CTIP, 2003b) outlined the development of two representative approaches to damage evaluation based on the emergent behaviour of the multi-agent system. The first of these employed the states of individual cells to collectively construct continuously-connected *impact boundaries* around damaged regions. This algorithm has been demonstrated on the Concept Demonstrator.

The second algorithm described in Report 4 develops *impact networks*, which are minimum spanning trees that connect a number of non-critical damage sites. This algorithm uses an Ant Colony System approach to the construction of the shortest paths between damage sites. It is one level more complex than the impact boundaries algorithm, because it connects groups of damaged cells whose members are within a predetermined distance of each other. This distance is set by the parameters of the ant algorithm.

Neither of these algorithms uses any information from the impact detection sensors, apart from the fact that a damage event has been detected. Damage is defined either by the failure, or partial failure, of intercellular communications, or by the simple detection of damage.

Work has continued in this area on a number of fronts during the past six months. Two areas in which significant progress has been made have aimed to develop emergent behaviour in more complex circumstances. In one case this has been the development of clusters of damaged cells, where clusters form (and re-form when new damage events occur) on the basis of the sensor signals. Thus, there is an emergent formation of clusters of cells with similar damage levels. This work will be described in the next Report.

A second problem takes a step towards developing the information required for self-repair, namely the self-replication of damaged shapes based on the emergent impact boundaries. Progress with the development of an emergent self-replication algorithm will be outlined below.

Shape Replication

In this section we investigate a possible first step towards the self-repairing ability, focusing, in particular, on the need for a robust self-replication of multi-cellular shapes. This *shape replication* ability should cover both ‘standard’ and ‘non-standard’ shapes. In other words, we expect that not only a standard shape predefined by an available structural ‘blueprint’ can be produced when required, but also that any non-standard and unpredictable shape covering a damaged region can be dynamically replicated on demand. Importantly, the self-replication mechanism should allow us to combine ‘standard’ and ‘non-standard’ shapes if necessary. Repair actions, such as shape replication, might be progressing in the environment where further impacts are likely to occur, and therefore there is a need for *robust* shape replication algorithms.

We investigate a multi-cellular shape replication mechanism implemented in a sensing and communication network, motivated by its application in robust self-monitoring and self-repairing aerospace vehicles. In particular, we propose:

- a self-referential representation (a ‘genome’), enabling self-inspection and self-repair;
- an algorithm solving the problem for connected and disconnected shapes; and
- a robust algorithm that can recover from possible errors in the ‘genome’.

The presented mechanism can replicate combinations of predefined shapes and arbitrary shapes that self-organize in response to occurring damage.

It is desirable that an impact boundary, self-organizing in the presence of cell failures and connectivity disruptions, forms a continuously-connected closed circuit. This circuit may serve as a reliable communication pathway around the impact-surrounding region within which communications are compromised. Every cell on a continuously connected closed circuit must always have two and only two neighbour cells, designated as the circuit members (circuit-neighbours of this cell). Algorithms that produce such impact boundaries have been developed in our earlier work, as mentioned above.

However, a continuously-connected closed impact boundary also provides a *template* for repair of the enclosed damaged region, uniquely describing its shape. Both these functions of impact boundaries can be contrasted with the non-continuous ‘guard walls’ investigated by Durbeck and Macias (2002) that simply isolate faulty regions of the Cell Matrix, without connecting elements of a ‘guard wall’ in a circuit.

The shape replication algorithms developed for this project are based on the principles of multi-cellular organization, cellular differentiation, and cellular division as well, similar to the embryonics approach described by Sipper et al. (1997) and Mange et al. (2000).

A desired shape is encoded when an emergent impact boundary self-inspects, and stores the ‘genome’ in a ‘mother’ cell. The genome contains both data describing the boundary, and a program of how to interpret these data. The mother cell is then seeded in a new location outside the affected region. Executing its program initiates *cell-replication* in the directions encoded in the genome. Each cell-replication step involves copying of the genome (both data and the program) followed by differentiation of the data: an appropriate shift of certain coordinates. Newly produced cells are capable of cellular division, continuing the process until the encoded shape is constructed (Prokopenko and Wang, 2004).

There may be cases when some fragments of the genome are damaged due to errors in the copying process or processor/memory failures. An advanced robust algorithm is designed to recover from such errors. The main challenge is not only to recognise corrupted data, but also to avoid a replication that may produce an incorrect shape. The proposed solution involves (1) self-inspection of the genome within a cell, determining the endpoints of disconnected boundary fragments, and (2) self-repair of the genome within a cell, filling the data between the disconnected fragments. Although the repaired genome does not cover all the missing cells, it does not introduce any cells that were not in the original shape, exhibiting soundness but not completeness. In other words, the repaired boundary is always contained within the original shape. Importantly, there is a redundancy in the shape replication process: other cells that did not suffer any damage would successfully replicate the parts not encoded in the partially repaired genomes (Figure 15).




Figure 15: Left: an impact boundary self-organizing around the damaged area. Right: the replicated shape.

7. The Demonstration System

At the time of this Report, the Concept Demonstrator is as shown in Figure 19. It has 188 physical cells instead of 192, because 4 have been removed to create the ‘window’ shown in the Figure 19 and clearly observable in the network views of Section 5. As explained in Sections 1 and 4 above, only four of these cells contain a Data Acquisition board to enable sensing to be carried out.



Figure 19: A general view of the Concept Demonstrator system (left), taken on April 16, 2004. The hexagonal prism structure covered in aluminium skin panels is at right, with computer monitors for visualization, fault diagnosis, and sensor signal monitoring on the left. One panel (four cells) has been removed to allow visual inspection of the cells. The image on the right is a view of some of the agents and power system. The grey connecting ribbon cables are the communication links between agents.

The sensing cells are mounted on the face of the hexagonal prism opposite to that containing the window, and a pulsed Nd:YAG laser is used to generate ‘impacts’ on these cells. Both the impact boundaries and impact networks algorithms were implemented.

Because the laser pulses used were unable to inflict critical damage on the CD cells, for the purposes of the impact boundaries algorithm damage regions were simulated around detected impacts. When an impact was detected, a damage region of radius proportional to the magnitude of the sensor signals was simulated by the inclusion of bit errors into the communication packets being sent by the ‘damaged’ cells. These errors are detected by neighbouring cells, which then recognise the damaged nature of the transmitting cells. An impact boundary is then formed around the damaged region.

Figure 20 contains visualizer screen views following two separate impacts of different laser pulse energy. A low energy pulse produced the small ‘damage’ region shown in the left screen, with a higher energy pulse producing greater ‘damage’ in the right-hand screen.

Figure 21 shows an expanded view of the impact region, showing the location of the impact as detected by the sensor arrays of the impacted and neighbouring cells. The impact has been detected by two adjacent cells, each of which calculates an impact location and positional uncertainty as outlined in Section 3.6. These detected locations are reported separately at this

stage, but later multiple detections of the same impact will be used to determine the most likely impact location.

In order to demonstrate the ability of the system to detect multiple impacts, and to operate effectively in the presence of damage, a cell in the sensing region was disabled and multiple impacts produced. A screen view from the visualizer, Figure 22, shows the results.

Figure 23 shows the demonstration system in its present form, with the visualizer screen next to the demonstration hardware.

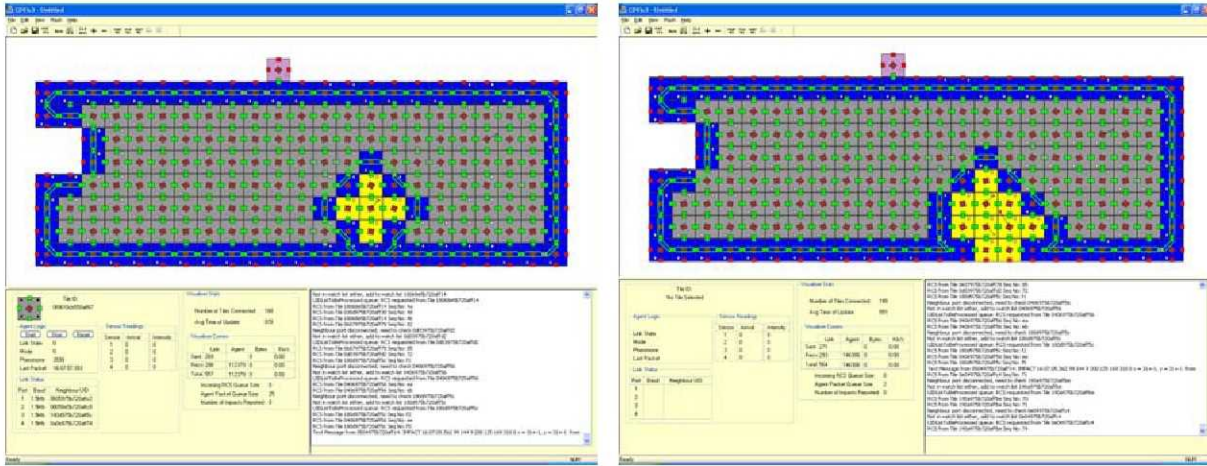


Figure 20: Boundary formation of a single sensed impact causing data corruption of radius one cell (left) and, after reassembly, for a higher energy impact (right).




Figure 21: Location of an impact detected by two adjacent cells. These were the only two cells on the panel that had sensing capability. Each cell provides an estimate of the impact location relative to the cell centre, and an estimate of the location uncertainty. The purple rectangle is the estimate of the impact location made by the cell on the left: the area of this box is large to indicate the large positional uncertainty that is associated with impact detection remote from the sensing cell (see Appendix A). The small square region (within the larger purple box) is the location estimated by the cell on the right: the positional uncertainty in this case is small because the impact is on the sensing cell.

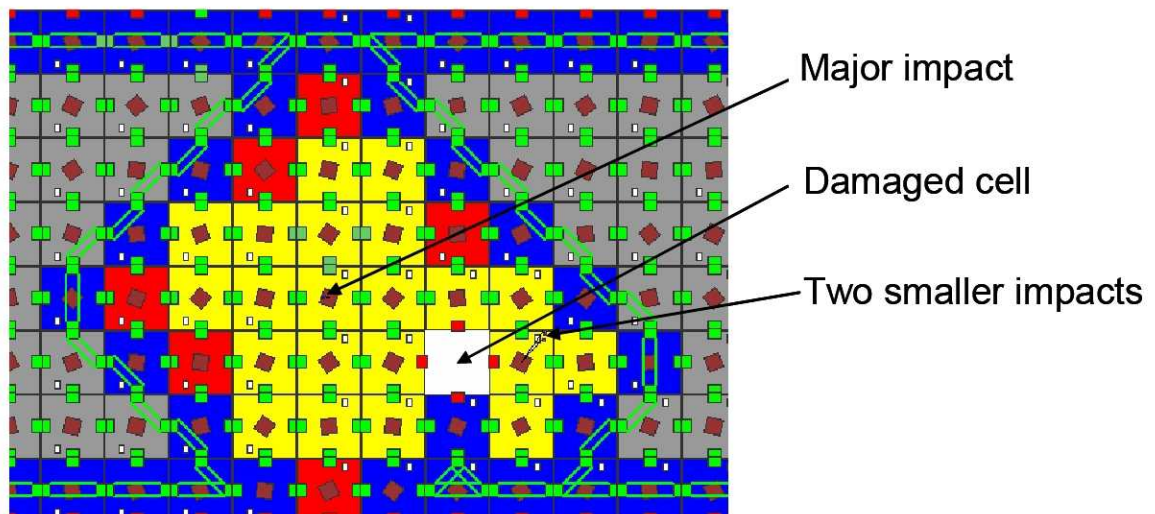


Figure 22: A screen view showing the boundary formed around the damage caused by three impacts, in a region containing a disabled cell. A high-energy impact was simulated on the indicated cell, and two low-energy laser impacts, almost coincident with each other, were detected as shown in the Figure.




Figure 23: The demonstration system following a laser pulse impact. The resulting impact boundary is shown on the visualizer screen.

8. Summary and Conclusions

This Report describes the second phase of the implementation of the Concept Demonstrator/Experimental Test-bed system containing sensors and processing hardware distributed throughout the structure, which uses multi-agent algorithms to characterize impacts and determine a suitable response to these impacts. It expands and adds to the report of the first phase implementation (CTIP, 2003b).

The current status of the system hardware is that all 192 physical cells (32 on each of the 6 hexagonal prism faces) have been constructed, although only four of these presently contain data-acquisition sub-modules to allow them to acquire sensor data. These cells incorporate a second version of the NAS electronics, which has increased flexibility through the use of an FPGA to implement the NAL communications, and improved functionality. The first production version of the DAS electronics has been designed, de-bugged and two samples manufactured, but a number of delays and manufacturing problems by outside suppliers and manufacturers have prevented the inclusion of more than a token number of boards at this time. Nevertheless, impact detection, location and severity have been successfully demonstrated. The software modules for simulating cells and controlling the test-bed are fully operational, although additional functionality will be added over time. The visualization workstation can now display additional diagnostic information about the array of cells (both real and simulated) and additional damage information.

Local agent algorithms have been developed that demonstrate emergent behaviour of the complex multi-agent system, through the formation of impact damage boundaries and impact networks. The system has been shown to operate well for multiple impacts, and to demonstrate robust reconfiguration in the presence of damage to numbers of cells. The next phase of the project will investigate further characterization of damage by implementing a level of secondary sensing, and will further develop the ability of the multi-agent system towards damage diagnosis, prognosis and an emergent response. This system is a powerful and flexible experimental platform for developing sensors, communications systems and multi-agent based algorithms for an intelligent vehicle health monitoring system for deployment in space vehicles.

9. References

- CTIP (2001). *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles*, Report 1. CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1516, 2001. Also published as NASA Technical Report no. NASA/CR-2002-211772.
- CTIP (2002). *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Development of Concepts for an Intelligent Sensing System*. Report 2. CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1517. May 2002. Also published as NASA technical report NASA/CR-2002-211773.
- CTIP (2003a). *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 3: Design of the Concept Demonstrator*. CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1628. January 2003. Also NASA/CR-2008-215307.
- CTIP (2003b). *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 4: Phase 1 Implementation of the Concept Demonstrator*. CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1898. Sept. 2003. Also published as NASA/CR-2008-215542.
- Durbeck, L.J.K. and Macias, N.J. (2002). *Defect-tolerant, fine-grained parallel testing of a Cell Matrix*. In Schewel, J. and James-Roxby, P.B. and Schmit, H.H. and McHenry, J.T. (Eds.) Proceedings of SPIE ITCom 2002 Series, Vol. 4867, 2002.
- Hedley, M., Johnson, M., Lewis, C., Carpenter, D., Lovatt, H., Price, D. C., 2003. *Smart Sensor Network for Space Vehicle Monitoring*. In Proceedings of the International Signal Processing Conference, Dallas, Texas, March 2003.
- Ilyas, M., (2003). *Ad Hoc Wireless Networks*, CRC Press, Florida, USA.
- Mange, D., Sipper, M., Stauffer, A. and Tempesti, G. (2000). *Towards robust integrated circuits: The embryonics approach*. In Proceedings of the IEEE, 2000, vol. 88, pp. 516-541.
- Prokopenko, M. and Wang, P. (2004). *On Self-referential Shape Replication in Robust Aerospace Vehicles*. To appear in Proceedings of the International Conference on the Simulation and Synthesis of Living Systems, ALIFE9, September 2004, Boston, USA.
- Simpson, W. (1994a). *The Point-to-Point Protocol (PPP)*, RFC1661, <http://www.faqs.org/rfcs/rfc1661.html>.
- Simpson, W. (1994b). *PPP in HDLC-like Framing*, RFC1662, <http://www.faqs.org/rfcs/rfc1662.html>.
- Sipper, M., Mange, D. and Stauffer, A. (1997). *Ontogenetic hardware*. BioSystems, vol. 44, pp. 193-207.
- Tanenbaum (2003). A. S., *Computer Networks*, Prentice Hall PTR, New Jersey, USA.

Appendix A: Location of Impacts

A.1 Introduction

This Appendix is a report on the performance of an experimental impact location system. The system consists of an array of four piezo-electric sensors located at the corners of a 60 mm square on an aluminium sheet substrate. The sensors are connected to a processor, which detects the arrival time of the elastic wave generated by the impact. This analysis takes no explicit account of dispersion of the elastic waves propagating in the aluminium sheet: it tacitly assumes that the leading edges of the impact-induced elastic disturbances detected by all four sensors contain the same combinations of frequencies. This is expected to be a reasonable assumption as long as the low frequency (leading) edge of the propagating S_0 Lamb wave is detected, since this mode is almost dispersionless at low frequencies, and/or when propagation distances are not too long (the precise meaning of this will be examined later on).

The experimental data used in this analysis was derived from a number of impacts that were simulated by short, high-power laser pulses focused onto an aluminium sheet in the vicinity of a group of four sensors (see Section 3 above and Report 4 [CTIP, 2003b]). The data consisted of the detected times of arrival of the impact-induced elastic disturbance at the four sensors, for impacts at a number of locations relative to the sensors. Figure A1 shows the positions of the sensors (red) and the impact locations (blue). The circle shown encloses the sensors and is of radius $60/\sqrt{2}$ mm.




Figure A1: Geometry of the sensors (red) and the impact locations (blue) relative to the sensors.

The four sensors are connected to a processor, which samples the data at a rate of ~ 3.125 MSPS. When a signal produced by a sensor is above a specified noise threshold, the time of arrival of the elastic wave is recorded. The speed of propagation is about 1.6 mm per clock count (equivalent to 5000 m/s). [This assumed propagation velocity is too low, the correct value being more like 5300 m/s for the low frequency components of the S_0 (extensional)

Lamb wave, but this difference is not expected to have a significant effect on the qualitative nature of the results presented below]. From these data it is required to determine the impact location relative to the sensors. The signals from the four sensors indicate the wave arrival times relative to the local time in the processor. As the time of the impact is not known, the problem is to determine both the impact time (relative to the local clock) and the $\{x, y\}$ position of the impact.

The above description is equivalent to the problem of radio-location such as that used by GPS, and thus the same mathematical techniques can be applied. These techniques have been used in other CSIRO projects. In summary, two methods of solving the problem are appropriate, namely a least squares fitting technique, and an analytical technique. Details of performance of these two methods are given below.

A.2 Analysis

This section describes the position fixing performance of both the least squares fitting algorithm and the analytical position determination algorithm. The data set consists of 24 impacts with the associated four arrival times measured in integral units of clock ticks (320 ns). As the speed of propagation is about 1.6 mm per clock tick, the accuracy of the system cannot be greater than about ± 1 mm. However, geometrical factors also affect the positional accuracy, and these effects are considered in the following sub-section.

Geometric Dilution of Precision

The accuracy of the position fix depends both on the time-of-arrival measurement accuracy and also on geometrical factors. The geometric factor, called the Geometric Dilution of Precision (GDOP), has a significant effect when the impact point is outside the circle shown in Figure A1. Within the circle GDOP is about unity, but outside the circle GDOP increases rapidly. In particular the GDOP at the centre of the circle with N sensors can be shown to be given by:

$$\text{GDOP}_{\text{centre}} = \sqrt{\frac{\sigma_x^2 + \sigma_y^2}{\sigma_d^2}} = \frac{2}{\sqrt{N}} \quad (\text{A1})$$

where σ_d is the standard deviation in the pseudo-range measurement, and σ_x and σ_y are the standard deviations in two orthogonal directions (not necessarily those of the sensor axes). With four sensors, the GDOP at the centre is unity. Further, it can be shown that, for N sensors, GDOP anywhere on the circle shown in Figure A1 is a constant value given by:

$$\begin{aligned} \text{GDOP}(N) &= \frac{2}{N^2 - 2C_N} \sqrt{\frac{N^4 - 3C_N N^2 + 2C_N^2}{N}} \\ C_N &= 1 + S_N^2 \\ S_N &= \sum_{i=0}^{N-1} \sin\left(\frac{\pi i}{N}\right) \end{aligned} \quad (\text{A2})$$

With four sensors the GDOP on the circle passing through the sensors is thus 1.978.

A particular case of interest is that of the radial line passing through the sensor location, as shown in Figure A2. (Note that the axes in Figure A2 are different from those used for the measurements in Figure A1).




Figure A2: Geometry of the radial line passing through the location of a sensor. The impact point is assumed to be outside the circle enclosing the sensors.

Inside the circle, GDOP can be computed to be:

$$\text{GDOP} = \sqrt{\frac{(3 + s^2)}{2s^2(3 - s^2)}} \quad (\text{A3})$$

$$s = \sin \theta = \frac{L}{\sqrt{2} r}$$

At the centre, GDOP is unity, in agreement with equation (A1). On the inside edge of the circle (at $\theta = \pi/4$), GDOP can be calculated to be $\sqrt{7/5} = 1.183$. There is a sharp transition to the value 1.987 on the circle, as derived from equation (A2) (see above).

Outside the circle the characteristics of GDOP similarly can be analysed analytically, resulting in the expression:

$$\text{GDOP} = \frac{\sqrt{2(c+1)(c+3)}}{2s^2} \quad (\text{A4})$$

$$s = \sin \theta = \frac{L}{\sqrt{2} r} \quad c = \cos \theta$$

Thus, GDOP exhibits an interesting behaviour as a function of the radial distance. Just on the outside of the circle the GDOP has a step increase to 3.588 according to equation (A4). The value of GDOP as a function of the radial distance is shown plotted in Figure A3.

Thus, GDOP has the curious feature of being triple-valued near the circle at a sensor point, one value (1.183) on the inner edge, a second value (1.987) on the circle, and a third value (3.558) on the outside of the circle. It is discontinuous for ideal point sensors, but not for real sensors of a finite size. Such discontinuities do not occur at other positions away from the sensor point.




Figure A3: GDOP calculated along a radial line that passes through a sensor at about 42 mm from the centre.

At greater radial distances, GDOP rapidly increases. At large distances from the origin ($s \rightarrow 0, c \rightarrow 1$), equation (A4) shows that GDOP approaches the function:

$$\text{GDOP} \approx \frac{2}{s^2} = \left(\frac{r}{L} \right)^2 \quad (\text{A5})$$

Thus at large distances from the origin (outside the enclosing circle), GDOP increases as the square of the radial distance measured in units of the square size. Because of the square law behaviour of the GDOP, the accuracy of the position fix rapidly decreases outside the enclosing circle. This behaviour is evident in the computed positions derived from the measured data, as will be demonstrated in the following subsection (see also Figure 8(b) in Section 4).

The large value of GDOP at large radial distances can be understood by the following qualitative argument. When the impact point is remote from the sensors a small change in the radial position (say by Δr) will result in the distances to all the sensors also increasing by about Δr . However, the fitting process uses pseudo-ranges rather than ranges, so that the common increment Δr is interpreted as a change in the impact time offset parameter, rather than a change in the ranges. Thus the pseudo-ranges of the original point and the incremental point are essentially the same, resulting in poor position determination accuracy.

The GDOP over the square area from -100 mm to +100 mm is shown plotted in Figure A4. The rapid increase in GDOP outside the central region is evident.

Least squares Fitting Algorithm Performance

This section describes the performance of the least squares fitting algorithm for determination of impact location from the experimental time of arrival data. Details of the algorithm are beyond the scope of this Report.

The least squares fitting algorithm is based on the approximate linearisation of the fitting equations, and thus requires many iterations to converge to a solution of sufficient accuracy.




Figure A4: GDOP plotted over the square area from -100 mm to +100 mm. The sensors are located at ± 30 mm from the centre of the figure in both vertical and horizontal directions.

For this Report the convergence limit was set at 0.1 mm, with the maximum number of iterations set at 10. The initial position was always set at the origin of the (x, y) coordinate system. The number of iterations required is shown in Figure A5. As can be observed the number of iterations required varied from 3 to a limiting maximum of 10 iterations, but typically the number of iterations did not exceed 6.




Figure A5: Number of iterations for determining the fit. The maximum is set at 10 iterations.

The LSF algorithm also requires that each measurement is given a weighting that represents the ‘quality’ of the measurement. For the analysis in this Report it is assumed that all measurements have equal weighting. Normal practice is to weight the data by the signal-to-noise ratio associated with each measurement. The SNR was not available, but the signal amplitude was measured. Thus assuming the noise is constant for each measurement, the SNR would be proportional to the signal amplitude. However, this enhancement was not used for the results presented in this Report.

The quality of the positional determination can be estimated by comparing the measured pseudo-range data with the pseudo-range data calculated from the computed position fix. The RMS error averaged over the four sensors is shown in Figure A6. If the RMS fit error is small, the data agrees well with the assumed mathematical model, while a large error probably indicates one or more measurement errors in the arrival times. Note that the RMS error is not the positional error, although the fit error and the positional error are usually highly correlated. As can be observed in the figure, the fit error is typically less than 1 millimetre (in agreement with the expectations discussed in the Introduction). However, some of the impacts (#0, #6 and #12) have much larger errors, and consequently the deduced positions of these impacts are expected to be less accurate.




Figure A6: Computed RMS fit error.

The computed positions of the impacts are shown in Figure A7. This diagram should be compared with the data in Figure 1. The positional errors are shown in Figure A8. As can be observed, many of the positional errors cluster about the origin (many points are overlaid in the Figure A8 and thus cannot be distinguished), but there are some very large errors of the order of 100 mm.




Figure A7: LSF computed positions of the impact.




Figure A8: Errors in the computed positions of the impact.

The distribution of positional errors can be more clearly observed from the data in Figure A9. In this figure, the magnitudes of the errors are sorted in ascending order. The data in the figure appear to be grouped into three categories. The first category contains those with small positional errors (say less than 4 mm, impacts #0 to #10). The second group (#11 to #22) have errors in the range 5 to 30 mm; these errors are presumably due to geometric factors (GDOP). Finally two impacts have a large error, probably due to a combination of GDOP and measurement errors.




Figure A9: Errors in the computed positions of the impact. The data are sorted in order of increasing error, so the impact numbers are different from those in other figures.

The nature of the errors can be further appreciated by plotting the errors as a function of radial distance from the coordinate origin. Figure A10 shows that the errors are approximately constant (about 2 mm) up to the circle radius, and then the errors increase rapidly as the radial distance increases. This behaviour is in accordance with the expectations of the analytical GDOP theory described in above.




Figure A10: The positional errors plotted as a function of radial distance of the impact. It can be observed that the errors are approximately constant up to the radial distance equal to that of the circle enclosing the sensors ($R_s \sim 42$ mm), and thereafter increase rapidly with radial distance.

This characteristic is further illustrated by plotting the errors as a function of GDOP, as shown in Figure A11. As can be observed, the errors are approximately proportional to GDOP, as expected from the theory. From the data in Figure A11, the positional accuracy is given approximately by:

$$\varepsilon = 2.5 * \text{GDOP mm} \quad (\text{A6})$$




Figure A11: The position errors are plotted as a function of GDOP. It can be observed that the errors are approximately proportional to GDOP.

Analytical Fitting Algorithm Performance

An alternative method of determining the position from the measured time-of-arrival data is by an analytical method originally proposed by Bancroft in 1985 for the GPS. This method can be adapted to the two-dimensional situation of this application. However, there is one complicating factor that is not relevant to the GPS case. For accurate position fixes the measured data must be suitably weighted. In particular, the optimum weighting is by the inverse of the distance from the impact point to the sensor. As the problem is to determine the position of the impact, the analytical solution must be performed twice. The first invocation of the algorithm assumes equal weighting for all the measured data, thus obtaining a first estimate of the position. Using this position the measurement weightings are calculated, and the algorithm invoked for a second time. Thus the analytical solution is also iterative, so there may be little benefit from using the analytical method.

The only situation when there may be some benefit is when the data are of poor accuracy. Figure A12 shows the data for the analytical and least-squares method for the impact data used in the previous section. As can be observed the positional estimates are essentially identical, except where the errors are largest (poorest data), where the analytical fit is superior. Thus there may be a case for the use of the analytical method when the LSF method has large fit errors. However, by comparing the fit errors with those in Figure A6, it can be observed that the worst cases of fitting (impacts #0, #6 and #12) do not correspond to points of failure of the LSF method.

The positional accuracy plotted against GDOP in Figure A13. As expected, the data are highly correlated, again showing that GDOP is the major factor in the determination of positional accuracy.




Figure A12: Comparison of the x-coordinates generated by least-square fit (LSF) and analytical methods. Note that in general there is good agreement between the two methods.




Figure A13: Positional error (analytical method) plotted as a function of GDOP.

A.3 Conclusions

Thus, in conclusion, the accuracy of the position fix within the circle surrounding the sensors is approximately constant, with the error determined by the time of arrival measurement accuracy. Outside the circle the positional errors increase rapidly, with the GDOP factor being the dominant effect. Indeed, for positional accuracy of (say) a few millimetres, the impact point should be within circles surrounded by sensors. Thus the impact measurement system must utilise the data from surrounding sensor arrays to ensure good positional accuracy.

Appendix B: Publications, Reports, Conference Papers and Articles Submitted

This Appendix lists the research output of the project group and collaborators since the beginning of the work in May 2001.

1. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles*, Report 1. D. Abbott, S. Cunningham, G. Daniels, B. Doyle, J. Dunlop, D. Economou, A. Farmer, D. Farrant, C. Foley, B. Fox, M. Hedley, J. Herrmann, C. Jacka, G. James, M. Johnson, B. Martin, G. Poulton, D. Price, T. Reda, G. Rosolen, A. Scott, P. Valencia, D. Ward, J. Winter, A. Young. CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1516, 2001. Also published as NASA Technical Report no. NASA/CR-2002-211772.
2. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Development of Concepts for an Intelligent Sensing System*. Report 2. D. Abbott, B. Doyle, J. Dunlop, A. Farmer, M. Hedley, J. Herrmann, G. James, M. Johnson, B. Joshi, G. Poulton, D. Price, M. Prokopenko, T. Reda, D. Rees, A. Scott, P. Valencia, D. Ward, J. Winter. CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1517. May 2002. Also published as NASA technical report NASA/CR-2002-211773.
3. *Concepts for an Integrated Vehicle Health Monitoring System*. D. Abbott, B. Doyle, J.B. Dunlop, A.J. Farmer, M. Hedley, J. Herrmann, G.C. James, M.E. Johnson, B. Joshi, G.T. Poulton, D.C. Price, M. Prokopenko, T. Reda, D.E. Rees, D.A. Scott, P. Valencia, D. Ward and J.G. Winter, Review of Progress in Quantitative Nondestructive Evaluation, Vol. 22, pp.1606-14 (eds. D.O. Thompson and D.E. Chimenti), American Institute of Physics Conference Proceedings Vol. 657, (2003).
4. *Experimental Observation of Linear and Non-Linear Guided Wave Propagation in Rolled Aluminum Sheets*. D.A. Scott and D.C. Price, Review of Progress in Quantitative Nondestructive Evaluation, Vol. 22, pp.1559-66 (eds. D.O. Thompson and D.E. Chimenti), American Institute of Physics Conference Proceedings Vol. 657, (2003).
5. *Theoretical Aspects of Linear and Non-Linear Guided Wave Propagation in Rolled Aluminum Sheets*. D.C. Price and D.A. Scott, Review of Progress in Quantitative Nondestructive Evaluation, Vol. 22, pp.1567-74 (eds. D.O. Thompson and D.E. Chimenti), American Institute of Physics Conference Proceedings Vol. 657, (2003).
6. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report 3: Design of the Concept Demonstrator*.

- D. Abbott, J. Ables, A. Batten, D. C. Carpenter, A. F. Collings, B. Doyle, J. B. Dunlop, G. C. Edwards, A. J. Farmer, B. Gaffney, M. Hedley, P. Isaacs, M. E. Johnson, B. Joshi, C. J. Lewis, G. T. Poulton, D. C. Price, M. Prokopenko, T. Reda, D. E. Rees, D. A. Scott, S. Seneviratne, P. Valencia, P. Wang, D. F. Whitnall, J. Winter.
CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1628.
January 2003. Also published as NASA Technical Report no. NASA/CR-2008-215307.
7. *Smart Sensor Network for Space Vehicle Monitoring.*
M. Hedley, M.E. Johnson, C.J. Lewis, D.A. Carpenter, H. Lovatt, D.C. Price,
In Proceedings of the International Signal Processing Conference (Dallas, Tx.), March 2003.
 8. *Self-reconfigurable Sensor Networks in Ageless Aerospace Vehicles.*
P. Wang, P. Valencia, M. Prokopenko, D.C. Price, G.T. Poulton,
Proceedings of the 11th International Conference on Advanced Robotics (ICAR-03),
Portugal, July 2003.
 9. *Self-organizing Impact Boundaries in Ageless Aerospace Vehicles.*
H. Lovatt, G.T. Poulton, D.C. Price, M. Prokopenko, P. Valencia, P. Wang,
Proceedings of the 2nd International Conference on Autonomous Agents and Multi-Agent
Systems (AAMAS-2003), Melbourne, July 2003.
 10. *Phase Transitions in Self-organising Sensor Networks.*
M. Foreman, M. Prokopenko, P. Wang,
Proceedings of the 7th European Conference on Artificial Life (ECAL-03), pp. 781-791.
Dortmund, Germany, September 2003.
 11. *An Integrated Health Monitoring System for an Ageless Aerospace Vehicle.*
D.C. Price, D.A. Scott, G.C. Edwards, A. Batten, A.J. Farmer, M. Hedley, M.E. Johnson,
C.J. Lewis, G.T. Poulton, M. Prokopenko, P. Valencia, P. Wang.
In "Structural Health Monitoring 2003", Proceedings of the 4th International Workshop on
Structural Health Monitoring, Stanford, CA, September 2003 (ed. F-K. Chang) DEStech
Publications, 2003, pp.310-18.
 12. *Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles. Report
4: Phase 1 Implementation of the Concept Demonstrator.*
D. Abbott, A. Batten, D. C. Carpenter, J. B. Dunlop, G. C. Edwards, A. J. Farmer, B.
Gaffney, M. Hedley, N. Hoshke, P. Isaacs, M. E. Johnson, C. J. Lewis, A. Murdoch, G. T.
Poulton, D. C. Price, M. Prokopenko, D. E. Rees, D. A. Scott, S. Seneviratne, P. Valencia,
P. Wang, D. F. Whitnall.
CSIRO Telecommunications and Industrial Physics. Report No. TIPP 1898.
Sept. 2003. Also published as NASA Technical Report no. NASA/CR-2008-215542.
 13. *Evolvable Recovery Membranes in Self-Monitoring Aerospace Vehicles.*
P. Wang, M. Prokopenko.
International Conference on Simulation of Adaptive Behaviour, Los Angeles, USA, July
2004. (accepted).
 14. *On Connectivity of Reconfigurable Impact Networks in Ageless Aerospace Vehicles.*
M. Prokopenko, P. Wang, M. Foreman, P. Valencia, D. C. Price, G. T. Poulton.

- Journal of Robotics and Autonomous Systems, in press, 2004.
15. *Self-organizing Hierarchies in Sensor and Communication Networks*.
M. Prokopenko, P. Wang, D. C. Price, P. Valencia, M. Foreman, A. J. Farmer.
Artificial Life Journal, Special Issue on Dynamic Hierarchies, 2004. (Accepted.)
 16. *On Self-referential Shape Replication in Robust Aerospace Vehicles*.
M. Prokopenko, P. Wang.
International Conference on the Simulation and Synthesis of Living Systems (ALIFE9),
Boston, USA, September 2004. (Accepted.)
 17. *Testbed for Structural Health Monitoring Network*.
M. Hedley, M. E. Johnson, C. J. Lewis, D. C. Price.
Global Signal Processing Conference and Expo, Santa Clara, USA, Sept. 2004. (Accepted.)
 18. *Communication Protocols for a Structural Health Monitoring Sensor Network*.
M. Hedley.
Paper submitted to 1st IEEE International Conference on Mobile Ad-hoc and Sensor
Systems Conference (MASS-2004), Florida, USA, October 2004.

| REPORT DOCUMENTATION PAGE | | | | Form Approved OMB No. 0704-0188 | |
|---|-------------|-------------------------------------|-------------------------------|---|---|
| <p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p> | | | | | |
| 1. REPORT DATE (DD-MM-YYYY) 01-09-2009 | | 2. REPORT TYPE Contractor Report | | 3. DATES COVERED (From - To) | |
| 4. TITLE AND SUBTITLE Development and Evaluation of Sensor Concepts for Ageless Aerospace Vehicles: Report 5 - Phase 2 Implementation of the Concept Demonstrator | | | | 5a. CONTRACT NUMBER | |
| | | | | 5b. GRANT NUMBER | |
| | | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) Batten, Adam; Dunlop, John; Edwards, Graeme; Farmer, Tony; Gaffney, Bruce; Hedley, Mark; Hoschke, Nigel; Isaacs, Peter; Johnson, Mark; Lewis, Chris; Murdoch, Alex; Poulton, Geoff; Price, Don; Prokopenko, Mikhail; Sharp, Ian; Scott, Andrew; Valencia, Philip; Wang, Peter; Whitnall, Denis | | | | 5d. PROJECT NUMBER PO L-71308D | |
| | | | | 5e. TASK NUMBER | |
| | | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NASA Langley Research Center Hampton, VA 23681-2199 | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER Report No. TIPP 2056 | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001 | | | | 10. SPONSOR/MONITOR'S ACRONYM(S) NASA | |
| | | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) NASA/CR-2009-215931 | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 38 Availability: NASA CASI (443) 757-5802 | | | | | |
| 13. SUPPLEMENTARY NOTES Langley Technical Monitor: Edward R. Generazio | | | | | |
| 14. ABSTRACT This report describes the second phase of the implementation of the Concept Demonstrator experimental test-bed system containing sensors and processing hardware distributed throughout the structure, which uses multi-agent algorithms to characterize impacts and determine a suitable response to these impacts. This report expands and adds to the report of the first phase implementation. The current status of the system hardware is that all 192 physical cells (32 on each of the 6 hexagonal prism faces) have been constructed, although only four of these presently contain data-acquisition sub-modules to allow them to acquire sensor data. Impact detection, location and severity have been successfully demonstrated. The software modules for simulating cells and controlling the test-bed are fully operational, although additional functionality will be added over time. The visualization workstation displays additional diagnostic information about the array of cells (both real and simulated) and additional damage information. Local agent algorithms have been developed that demonstrate emergent behavior of the complex multi-agent system, through the formation of impact damage boundaries and impact networks. The system has been shown to operate well for multiple impacts, and to demonstrate robust reconfiguration in the presence of damage to numbers of cells. | | | | | |
| 15. SUBJECT TERMS Nondestructive evaluation; NDE; Structural health monitoring; IVHM; NDI | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | STI Help Desk (email: help@sti.nasa.gov) |
| U | U | U | UU | 59 | 19b. TELEPHONE NUMBER (Include area code) (443) 757-5802 |